

12

Managing Risk with Chance-Constrained Programming

MICHAEL BEVERS AND BRIAN KENT

Reducing catastrophic fire risk is an important objective of many fuel treatment programs (Kent et al. 2003; Machlis et al. 2002; USDA/USDI 2001a). In practice, risk reductions can be accomplished by lowering the probability of a given loss to forest fires, the amount of probable loss, or both. Forest fire risk objectives are seldom quantified, however, making it difficult to compare alternative management options. Instead, fuel treatment programs are often selected based on expert judgment.

Relying solely on expert judgment can be effective for some decisions. For example, when treatment budgets are relatively small, fuel reductions in or fuel breaks around a few obvious high-risk areas may be the only affordable options. By considering ignition sources and probabilities, vegetation mosaics, terrain, and the property and resource values at risk, fire managers might easily make a number of good treatment choices without needing to compare quantified risk estimates. On the other hand, when fuel treatment budgets are substantially larger but too small to treat fuels everywhere, it might be helpful for managers to use mathematical models to analyze the complex interactions across a landscape. This approach seems particularly pertinent for informing decisionmakers about explicit management strategies and trade-offs when considering fires large enough to impact sizable portions of the landscape.

Fire managers have much experience with mathematical models that predict or simulate the spread, intensity, and resulting effects of particular wildfire events (e.g., Finney 1998). Many are also familiar with the use of mathematical programming models as tools for improving the efficiency of resource allocations,

such as optimally locating air tankers (Greulich and O'Regan 1982). In this chapter, we suggest the possibility of using mathematical programming, based in part on fire simulations, to improve the efficiency of fuel treatment programs by optimizing a quantified risk objective.

Determining fuel treatments for a forested landscape can be cast as a type of spatial optimization problem (Hof and Bevers 1998, 2002) where the effects of treatments in each area are taken into account not only in that area, but also in nearby areas to or from which fires might spread. To address risk, we suggest an optimization modeling method known as chance-constrained programming (e.g., Hof et al. 1992). We will first briefly introduce this type of chance-constrained programming problem and then explore what the approach might offer the fire management community.

Chance-Constrained Programming

Chance-constrained programming is one of the major mathematical modeling approaches for dealing with the effects of random events on management outcomes (Henrion 2005). Typical areas of application include engineering, finance, and resource management, where random variables such as traffic volume, product demand, market prices, weather, fire occurrence, and fire severity are inherent and important parts of the system under consideration. Chance-constrained programming models must cope with the difficulty of choosing the best decisions possible—the optimal ones given objectives—prior to observation of the random events affecting the outcomes that follow.

Consider the problem of selecting the best locations and types of fuel treatment for a landscape such as the forested area in the upper South Platte River basin of Colorado (see Figure 12-1). This area includes portions of the 138,000-acre Hayman fire, which burned in 2002, visible as a finger of nonforest and scattered tree cover along the lower left edge of the map. The map includes portions of other exceptionally large wildfires that have occurred in the past decade as well (Graham 2003). Given a limited budget for fuel treatment, fire managers here and on other fire-prone forests might want to choose treatments for the landscapes in their care that would minimize future losses to wildfires over the duration of time those treatments would remain effective. We can characterize this management problem mathematically.

Suppose the manager has divided the landscape in Figure 12-1 into 200 units, each with four treatment options. Considering the random nature of fire, many different scenarios of how future fires will interact with landscape treatments are possible. For each of these scenarios, the optimization problem is as follows:




FIGURE 12-1. Forest Fuel Types in the Upper South Platte River Basin of Colorado.

$$\text{Minimize } \sum_{i=1, \dots, 200} Z_i(t) \quad (12-1)$$

$$\text{subject to } \sum_{i=1, \dots, 200} \sum_{j=1, \dots, 4} c_{ij} X_{ij} \leq C \quad \text{with } X_{ij} \text{ either zero or one} \quad (12-2)$$

$$\sum_{j=1, \dots, 4} X_{ij} \leq 1 \quad \text{for } i = 1, \dots, 200 \quad (12-3)$$

$$Z_i(t) = f_i(\vec{X}, \vec{Y}, t) \quad \text{for } i = 1, \dots, 200 \quad (12-4)$$

where Equation 12-1 minimizes total wildfire losses, summed across all 200 possible treatment areas on the landscape. $Z_i(t)$ is a random variable accounting for loss to wildfires in treatment area i ; random losses, such as fire suppression cost and property damage, are unknown until treatments have occurred and the subsequently observed (or simulated) losses are later tallied at time t for this scenario. Equation 12-2 constrains total expenditures on fuel treatments to be less than or equal to C , the treatment budget. X_{ij} is a binary variable whose value indicates selection ($X_{ij} = 1$) or nonselection ($X_{ij} = 0$) of treatment program j (possibly specifying a variety of treatments) from the set of four optional treatment programs for each area i ; having all four treatment variables for an area set to 0 indicates that no treatments are planned. The cost of implementing any treatment program X_{ij} is entered as parameter c_{ij} . Equation 12-3 restricts the allocation in each treatment area to one program.

Equation 12-4 expresses the complexity of modeling wildfire losses for one scenario on a landscape subdivided into treatment areas, indicating that losses in each area are a function of the vectors $\vec{X}$, comprising the fuel treatment chosen for that area as well as, at least potentially, the treatments chosen for all other areas. Likewise, the vector $\vec{Y}$ of initial fuel conditions, resource values, property values, and other factors across the landscape potentially influences fire losses in area i , as do many random factors, such as weather and fire ignitions, expressed in equation 12-4 simply as a function of time t . The f_i function for each treatment area in equation 12-4 would have to account for many possible wildfires, each as a random outcome based on t , including fires originating in that area, surface or crown fires spreading into that area from adjacent areas, and perhaps the possibility of spot fires caused by wildfires burning in areas farther away, depending on the sizes of the treatment areas.

Modeling only one of many possible futures is a serious limitation, however. A set of fuel treatments that minimizes random losses for one future scenario will not necessarily work well for all possible futures. It is necessary to account for many possible futures in order to manage probable losses.

One commonly used approach is to optimize based on expected values of random variables. Mathematically, this would entail replacing equation 12-1 with

$$\text{Minimize } E \left[\sum_{i=1, \dots, 200} Z_i(t) \right] \quad (12-5)$$

to minimize expected loss and expanding equation 12-4 to estimate all possible $Z_i(t)$ and their probabilities of occurring in order to compute the

expectation (E). But this approach, usually regarded as risk-neutral (see Hardaker et al. 1997), seldom supports high reliability; that is, results are not low-risk (Hof et al. 1995). For example, suppose solving the problem using equation 12-5 resulted in \$10 million of expected loss (assuming for the moment that the losses of interest are purely financial). If modeled losses correctly reflected reality and were normally distributed, then a manager implementing the selected fuel treatments anticipating that future fire losses will not exceed \$10 million would have only a 50:50 chance of success versus failure. When considering severe fire losses, many managers might ask what the level of loss would be for 80:20, 90:10, 95:5, or even 99:1 odds, and what they could do to minimize those losses.

Chance-constrained programming is designed to answer these types of questions by allowing the manager to specify an acceptable level of risk. Here risk is defined as the probability of an unwanted outcome occurring. In the example above, risk could be the probability of random losses exceeding the level minimized in the mathematical program. In other literature, the term *reliability* or *safety* is sometimes used in place of risk; these terms refer to the probability of losses being at or below the planned (minimized) level (reliability = 1—risk). Although some may be uncomfortable with specifying an acceptable level of risk, costs associated with plans to handle extremely rare events tend to be unattractive enough that managers usually must accept some level of risk.

In choosing odds, the manager sets an acceptable level of risk, keeping in mind that more than one level might be modeled in a series of "what if" exercises to better inform the decisionmaker. A level of risk deemed to be acceptably low, but not too extreme, is entered into the model as a probabilistic or chance constraint. Mathematically, the objective function in the above model, previously equation 12-1 or 12-5, is replaced with

$$\text{Minimize } B \quad (12-6)$$

$$\text{subject to} \quad \Pr \left(\sum_{i=1, \dots, 200} Z_i(t) > B \right) < p \quad (12-7a)$$

The chance constraint, equation 12-7a, sets an upper bound B on total fire losses and requires that the probability (Pr) of exceeding those losses be less than parameter p , the acceptable level of risk (e.g., 0.05). The resulting probable loss B is minimized in equation 12-6. As with equation 12-5, this approach also requires that equation 12-4 be expanded to estimate random loss $Z_i(t)$ values and probabilities for many scenarios.

In practice, the same problem often is solved using an equivalent formulation based on acceptable reliability rather than acceptable risk. Equation 12-7a then becomes

$$\Pr \left(\sum_{i=1, \dots, 200} Z_i(t) \leq B \right) \geq q \quad (12-7b)$$

where probability parameter q is an acceptable level of reliability (e.g., 0.95). Many other related formulations are also possible (Hof and Bevers 1998).

Solving the chance-constrained programming problem presented by equations 12-2, 12-3, 12-6, 12-7a, and an expanded version of 12-4 for all possible combinations of fuel treatment budget (parameter C) and acceptable risk (parameter p), we would expect results that could be presented in three dimensions. First, parameter C from equation 12-2, expressed as a percentage, is varied from a zero budget allocation to full funding (100 percent). Parameter p from equation 12-7a, also expressed as a percentage, varies between no risk (or 100 percent reliability) to 100 percent risk (or no reliability). The third dimension presents the minimized loss B from equation 12-6. To offer a sense of scale for the percentages, a 100 percent loss represents the worst loss possible when no fuels are treated on the landscape. There is no risk that this level of loss will be exceeded; reliability is 100 percent that losses will not exceed this amount. A 100 percent budget represents the (unlikely) budget that would be required to make the landscape completely safe from fires, with no risk of any loss occurring.

The production surface defined over these three dimensions shows the smallest probable loss to fires that can be achieved given any combination of budget and risk (the probability that losses will be greater than indicated by this surface). Whether the optimal surface for an actual problem is concave, convex, or more complex would depend on the circumstances being modeled. Because each point on this surface is efficient, with loss minimized, trade-offs among budget, risk, and loss are revealed, providing the manager with useful information about reasonable levels of funding and risk to be considered.

It is not uncommon for trade-offs to indicate that large reductions in risk can be achieved, up to a point, with modest increases in expenditures. In taking advantage of such opportunities, managers need to recognize that minimizing probable loss for any combination of budget and risk requires that the budget be spent on a particular set of fuel treatments across the landscape, one that is optimal for (and potentially unique to) that combination. The mathematical solution to the chance-constrained programming problem specifies the necessary treatments, by location, as well as the resulting probable loss. As with any model, the results are best used to inform decisionmakers rather than being viewed as decisions. To that end, a model that could help identify which fuel treatments to use and where to apply them so as to minimize probable loss is potentially a powerful tool.

Once the manager has a firm fuel treatment budget, he or she can focus on a single line on the production surface. This line indicates trade-offs of interest to the manager between risk and loss for the given budget. When a risk parameter value is set in equation 12-7a, that constraint can also be drawn as a line on the production surface. Only points on or above both constraints are then candidates for selection; the point with the smallest loss in this constrained region of the surface is selected as the optimal solution (B^*) to this chance-constrained

programming problem. The next section illustrates these ideas with a hypothetical problem that is small enough to solve with manual calculations.

Applying the Concept in Community Fire Management

Let's look at a simple example. Suppose two rural communities, A and B, occupy the local forest. Both have already done substantial work to establish defensible space around structures, and fire managers now want to evaluate the benefits of putting additional fuel treatments around one or both communities (e.g., Agee et al. 2000; Bevers et al. 2004; Finney 2001; Lindenmayer and McCarthy 2002). Estimates of the resulting projected random losses in A given a large fire are distributed as $\text{Pr}(\text{loss} = 8, 9, 10, 11, 12 \text{ million dollars}) = (0.20, 0.20, 0.20, 0.20, 0.20)$; projected random losses in B given a large fire are distributed as $\text{Pr}(\text{loss} = 1, 2, 3, 4, 15 \text{ million dollars}) = (0.20, 0.20, 0.20, 0.20, 0.20)$. Losses would be greater in A than in B, with the exception of a high-value electronic site, for which the chance of loss is 20 percent during a large fire despite special protection plans.

While fuel treatments remain effective, the probability that a large fire starts only in A is 0.05, that a large fire starts only in B is 0.10, and that large fires start independently in both A and B is 0.005; the probability that no large fires start during that period is 0.845. Considered individually, the expected loss is about the same for each community, because the probability of a large fire is almost twice as high in B as in A. Some probability also exists, however, that a large fire starting in one community will spread to and severely burn the other; these conditional fire spread probabilities depend on the presence or absence of the additional fuel treatments. To simplify comparisons, we will assume that the cost for additional fuel treatments around A is the same as the cost for additional fuel treatments around B.

With no additional fuel treatments, the conditional probability that a large fire spreads is 0.40 from A to B and 0.30 from B to A. With treatments only around A, the probability of spread is 0.20 from A to B and 0.10 from B to A. With treatments only around B, the probability of spread is 0.13 from A to B and 0.15 from B to A. With treatments around both A and B, the probability of spread is 0.026 from A to B and 0.015 from B to A. Much of this information likely would come from simulations.

Given the information above, the expected loss with no additional treatments is \$1.475 million. With treatments only around A, the expected loss is \$1.225 million. With treatments only around B, the expected loss is \$1.2575 million. With both sets of treatments, the expected loss is \$1.0965 million. From an expected-value point of view, treatments around A are preferable to treatments around B if only one area can be treated. Ignoring alternative uses that might exist for the treatment budget, it would be worth treating fuels around community A as long as costs would not exceed \$0.25 million (the

reduction in expected loss); both sets of fuel treatments would be worthwhile as long as total costs would not exceed \$0.3785 million.

To consider the same problem from a chance-constrained perspective, we calculate the information presented in Tables 12-1 and 12-2. Table 12-1 shows the cumulative probabilities for large fires occurring at or below each level of loss under each of the four treatment options. Probabilities are reported to an insignificant number of decimal places considering the precision of our parameter set to show that the preferred treatment set can depend on the level of reliability desired. For losses ranging from \$12 million to \$14 million, at reliability levels from about 0.967 to 0.975, treating fuels only around B becomes a slightly better choice than treating fuels only around A; otherwise treatments around A are generally better. Table 12-2 uses the cumulative probability (rounded to three decimal places) for the optimal treatment at each level of loss to report efficient trade-offs for budget levels that would construct zero, one, or two sets of treatments.

TABLE 12-1. Cumulative Probabilities of Losses (Million\$) Being Less Than or Equal to Stated Amounts for Four Alternative Fuel Treatment Strategies

Loss	Cumulative probabilities			
	No treatments	Set A	Set B	Sets A and B
0	0.845	0.845	0.845	0.845
1	0.859	0.863	0.862	0.8647
2	0.873	0.881	0.879	0.8844
3	0.887	0.899	0.896	0.9041
4	0.901	0.917	0.913	0.9238
8	0.907	0.925	0.9217	0.93354
9	0.9152	0.934	0.93146	0.943592
10	0.9256	0.944	0.94228	0.953956
11	0.9382	0.955	0.95416	0.964632
12	0.953	0.967	0.9671	0.97562
13	0.9618	0.971	0.97134	0.976868
14	0.9684	0.974	0.97452	0.977804
15	0.9868	0.994	0.99364	0.998128
16	0.989	0.995	0.9947	0.99844
23	0.9912	0.996	0.99576	0.998752
24	0.9934	0.997	0.99682	0.999064
25	0.9956	0.998	0.99788	0.999376
26	0.9978	0.999	0.99894	0.999688
27	1.0	1.0	1.0	1.0

TABLE 12-2. Cumulative Probabilities of Losses (Million\$) Being Less Than or Equal to Stated Amounts under Optimal Management for Three Budget Levels

<i>Cumulative probabilities^a</i>			
<i>Loss</i>	<i>No treatments</i>	<i>One set</i>	<i>Both sets</i>
0	0.845	0.845	0.845
1	0.859	0.863	0.865
2	0.873	0.881	0.884
3	0.887	0.899	0.904
4	0.901	0.917	0.924
8	0.907	0.925	0.934
9	0.915	0.934	0.944
10	0.926	0.944	0.954
11	0.938	0.955	0.965
12	0.953	0.967	0.976
13	0.962	0.971	0.977
14	0.968	0.975	0.978
15	0.987	0.994	0.998
16	0.989	0.995	0.998
23	0.991	0.996	0.999
24	0.993	0.997	0.999
25	0.996	0.998	0.999
26	0.998	0.999	1.0
27	1.0	1.0	1.0

^aRounded to three decimal places.

Issues Associated with Chance-Constrained Programming

An obvious issue that arises with the use of chance-constrained programming is the technical challenge. Solving chance-constrained programming problems can be very difficult, and finding suitable mathematical algorithms remains an active research area. Formulating the details of equations 12-4 and 12-7a also requires substantial further research. Once sufficient progress is made on these two fronts, the question remains whether fire-planning staffs will be equipped for the mathematical programming effort required to implement these methods; real-time models are even further beyond current technology. Whether or not chance-constrained programming models are implemented for fire planning, one might still

gain insights into risk management issues by considering them from a risk trade-off point of view.

Multiple Objectives

Decision problems are often complicated by the need to address multiple objectives simultaneously. Earlier, we assumed for simplicity that all fire losses were financial and could be combined into a single measure of dollars for an objective function. The other objective, not to exceed the budget (also in dollars), could be handled as a deterministic (nonrandom) constraint. But real problems tend to be messier, sometimes posing a variety of objectives to be optimized that do not readily combine into a single measure such as dollars or a single acceptable level of risk. For example, a manager might decide that the prescribed-fire portion of a fuel treatment program warrants special constraints on smoke production and escaped fire costs. These, too, are random in nature and potentially could be modeled using chance constraints. The resulting changes to the chance-constrained objectives might look like this:

$$\text{Minimize } B_z + B_w \quad (12-8)$$

$$\text{subject to } \Pr \left(\sum_{i=1, \dots, 200} Z_i(t) > B_z \right) < p_z \quad (12-9)$$

$$\Pr \left(\sum_{i=1, \dots, 200} W_i(t) > B_w \right) < p_w \quad (12-10)$$

$$\Pr \left(\sum_{i=1, \dots, 200} V_i(t) > b_v \right) < p_v \quad (12-11)$$

where $Z_i(t)$ in equation 12-9 now accounts for random losses from fires other than escaped prescribed fires, $W_i(t)$ in equation 12-10 accounts for random losses from escaped prescribed fires, and $V_i(t)$ in equation 12-11 accounts for random smoke production that could concentrate in population centers within the airshed from prescribed fires. Each constraint has a separate risk parameter, and equation 12-11 has a fixed upper bound parameter b_v , reflecting an air quality standard to be met (as in Fuessle et al. 1987); it is not necessary to minimize smoke production given that standard. Although the acceptable level of risk for losses from escaped prescribed fires differs from that for other forest fires, the financial losses B_z and B_w are commensurable in the objective function, equation 12-8.

Different measures to be optimized are not always commensurable. Much mathematical programming literature (e.g., Cohon 1978; Steuer 1986) has been devoted to the problem of simultaneously addressing multiple objectives by formulating linear combinations using various weights as coefficients,

penalizing deviations from predetermined targets, or other methods that reflect the relative importance of each objective. An issue highlighted by chance-constrained programming is that decisionmakers need to consider not only the appropriate weight for each objective, but also the appropriate level of acceptable risk. As equations 12-8 through 12-11 illustrated, managers trying to meet multiple objectives will not necessarily be equally risk-averse to every objective. To offer another example, although private property damage, suppression cost, and watershed rehabilitation cost, all measured in dollars, might be weighted equally in an objective function, managers could be less willing to accept risk regarding the cost of damages to private property than for public expenditures such as fire suppression and watershed rehabilitation. Each objective handled by a chance constraint, whether included in the objective function or not, requires the specification of an acceptable risk of failure.

Reducing acceptable risk for one objective can result in poorer achievement for others, similar to the effect of increasing the weight in the objective function of one objective relative to others. Although some degree of interplay exists between importance and risk, making it more difficult to balance conflicting objectives for random systems, managers should avoid the temptation to address a need for high reliability by assigning greater importance to an objective. Balancing the importance of and acceptable risk for several objectives can be especially complicated when stakeholder groups vary by objective, as is often the case. Nevertheless, the effort should lead to better management of random systems.

Diversification and Hierarchical Differences

In bureaucratic organizations, typified by a hierarchy of authority, managers at different levels of the organization may have different acceptable levels of risk for objectives. This tends to occur even when those managers jointly agree on the relative importance of each objective.

Spatial diversification is a principal underlying cause for these differences in risk avoidance. Although a regional- or national-level manager is responsible in a supervisory capacity for local forest fire management, large fire losses at some locations tend to be offset by small losses at others. For the local manager, on the other hand, large fire losses, though infrequent, are more catastrophic. This problem is ameliorated to some extent by at least two factors. First, because catastrophic fire losses are somewhat unavoidable from a large-scale perspective, regional or national managers plan for these occurrences. Second, because large fire costs are extremely high and the number of large fires varies substantially from year to year, annual fire costs are quite variable even when accumulated nationally. So regional and national managers are not entirely buffered by diversification from the consequences of bad fire years. Nevertheless, some tendency is likely for

managers to be more risk-averse the lower they are organizationally. In the models discussed earlier, this might mean that a local manager would prefer for planners to use a smaller risk parameter (p) in equation 12-7a for some management objectives, whereas a higher-level manager might be comfortable using a larger risk parameter or perhaps even using expected values instead of chance constraints—in other words, equation 12-5 instead of equations 12-6 and 12-7a.

Bureaucratic policies and practices can also lead to differences in acceptable levels of risk at different levels of the organization. For example, suppose a regional manager has \$10 million to allocate this year as \$1 million budgets to each of 10 local area managers in an organization where managers are expected to completely spend their funds in order to support next year's budget request. If the regional manager is willing to take a 20 percent chance of overspending, each local manager can take only about a 2 percent chance (assuming independence for simplicity). Again, based on diversification, local managers might tend to be more risk-averse than regional or national managers.

Despite the advantages presented by spatial diversification, upper-level managers can be more risk-averse than local managers in some circumstances. This can result from simple differences in individual personalities, but other influences are also possible. For example, promotions to upper-level management positions might be based on patterns of risk avoidance that have become ingrained. Where promotions are based on successful achievements, failures of commission (mistaken actions) are sometimes more glaring than failures of omission (mistaken decisions not to act). In many cases, the information available to upper-level managers is filtered and summarized to the point where it lacks details that might support taking risky actions, or upper-level managers might not have the lower-level experience needed to understand fully the information at hand. Conversely, lower-level managers sometimes lack the information or experience to appreciate the broader context of risky decisions, including the political influences from outside organizations and elected officials. Although organizations function in complicated environments that may be designed in part to compensate for such effects, hierarchical differences in risk tolerance are likely.

Repeatability

Much like diversification, making decisions based on expected values relies on repeatability, unless one is indifferent to the risk of a poor outcome (Christiansen 1979); given repeated random chances, the average outcome is likely to fall close to the expected value. But repeatability, or the lack of it, can influence decisions regardless of one's level of risk aversion.

Let's reconsider our simple fuel treatment example from the standpoint of repeatability. A government agency, or even an insurance company, might

view fuel treatments around one or both communities as cost-effective based on expected values, especially if it will be conducting similar fuel treatments around dozens of other communities at scattered locations. On the other hand, if the communities have to raise the funds for treatments themselves, the information in Table 12-2 and the issue of nonrepeatability might cause them at least to question the investment and perhaps abandon it, for two reasons. First, only a 0.155 probability exists of any loss occurring. Second, the marginal reductions in risk resulting from fuel treatments are quite small. The largest possible reduction in risk from installing one set of treatments is only 0.019, at a random loss of \$15 million. From a local point of view, only one of the losses in Table 12-2 can occur over the planning time frame, offering no opportunity to accumulate multiple benefits from repeating many small risk reductions. In this example, the more detailed risk information in Table 12-2 might lead community-level decisionmakers (despite having to endure the consequences without benefit of diversification) to choose a riskier alternative when they have to pay for fuel treatments, because nonrepeatable treatment benefits are small. Matching grants could be used by governments in cases like this to encourage projects on an expected-value basis with local communities paying only for a nonrepeatable share of the benefits.

Conclusions

In this chapter, we considered the use of mathematical programming models to help people choose among various wildfire risk-reduction measures. We observed that, in contrast with government entities, which might choose to minimize total expected loss, residents of forest communities often face non-repeatable or infrequent chances of catastrophic loss to wildfires. This observation led us to examine chance-constrained programming as a possibly useful mathematical programming method.

Risk of loss to forest fires, as defined here, is unavoidable yet manageable. Indeed, risk is an integral factor between managerial effort and outcomes in random systems. As chance-constrained programming capability continues to evolve, fire managers might one day be better able to quantify risk and help identify preferred actions with these mathematical models. Chance-constrained models of fuel treatments are now tractable for small landscapes and might soon be for larger, more realistic cases.

Meanwhile, it should be useful for managers to recognize the greater scope of the fire-planning problems they face. Realizing that objectives ranked as similar in importance do not necessarily have similar acceptable levels of risk, which might vary not only across stakeholder groups, but also across organizational levels, should lead to more perceptive and creative approaches to problem solving. We hope this is the case while researchers continue to develop chance-constrained programming as a tool for fire management.

Acknowledgments

We are grateful to Merrill Kaufmann (USDA Forest Service) and Ingrid Burke and Jiong Jia (Colorado State University) for the fuel type map used in Figure 12-1. We also thank Bruce Meneghin and Jim Saveland (USDA Forest Service) for helpful suggestions, along with Wade Martin (California State University, Long Beach).