

Knowledge partitioning: Context-dependent use of expertise

STEPHAN LEWANDOWSKY and KIM KIRSNER

University of Western Australia, Nedlands, Western Australia, Australia

Although exceptional performance is a defining attribute of expertise, experts sometimes exhibit striking errors and performance limitations. This article reports two experiments in which experts predicted the spread of bush fires, a domain characterized by complex but well-understood physical dynamics. Although accuracy was typically high, large errors were observed when two primary predictor variables were in opposition. In a second study, the experts' behavior—in contrast to that of novices—was additionally shown to depend on problem context. In one context, experts again committed errors, whereas in another, equally domain-relevant context, the correct predictions were made. Critically, when comparing performance across contexts, completely opposing predictions were made under identical physical conditions. We therefore suggest that expertise may comprise separate, and sometimes even mutually exclusive, components of knowledge.

There is much support for the view that expertise, the ability to perform exceptionally well in a domain, arises primarily from extensive deliberate practice, rather than from innate ability or talent (e.g., Ericsson & Charness, 1994; Ericsson & Lehmann, 1996). The practice view entails several assumptions about the nature of expertise, two of which are particularly relevant here. First, expertise is highly domain specific and does not represent a general characteristic or trait (see, e.g., Ericsson & Lehman, 1996). Second, the specificity of expertise is compatible with the finding that experts are susceptible to systematic errors, biases, and limitations of performance (see, e.g., Camerer & Johnson, 1991; Sternberg, 1996; Sternberg & Frensch, 1992).

This article explores the limitations of expertise in the prediction and control of bush fires, a domain ideally suited for assessing performance, because the underlying physical processes are reasonably complex yet well under-

stood. Although accuracy was found to be high in most situations, systematic errors occurred on some domain-relevant problems when they were presented in one particular context. When the same problems were presented in another, equally domain-relevant context, experts made the correct prediction. The apparent context specificity of performance and the associated incompatibility of predictions across contexts suggested that expertise may comprise separate, and sometimes even mutually exclusive, components of knowledge.

EXPERTISE AND ITS LIMITATIONS

The specificity of expert knowledge is particularly evident when individuals practice a task in the laboratory over prolonged periods and reach outstanding performance levels. For example, deliberate practice of mnemonic techniques can raise digit span from the standard 7 ± 2 to 80 or even higher (see, e.g., Ericsson, Chase, & Faloon, 1980; Staszewski, 1993). Verbal protocols gathered during training revealed that subjects enhanced their digit spans by creating increasingly larger, richly integrated hierarchical retrieval structures (e.g., Staszewski, 1993). Notwithstanding that dramatic increase in digit span, the capacity to encode other information, remarkably, remained invariant at approximately seven symbols (Chase & Ericsson, 1981).

The view that expert knowledge is richly integrated, yet domain limited, is compatible with the fact that experts sometimes are neither outstanding nor better than novices. Two of these expert fallacies—inflexibility and expediency—are particularly relevant here. Inflexibility arises when experts cannot adapt to a change in the deep structure of a problem as readily as novices (Frensch & Sternberg, 1989; Sternberg & Frensch, 1992). For example, Sternberg and Frensch showed that expert bridge players suffered more than novices from a changed rule's

Preparation of this article was facilitated by Large Research Grant A79600016 from the Australian Research Council to the first author and by Industrial Collaborative Grant AM9280018 from the Australian Research Council to K.K., John Dunn, David Morrison, and C. P. Tsang, held jointly with Australian Defence Industries Ltd., Perth. This article was completed while the first author held a position as a Visiting Professor at the University of Potsdam, Germany, financed by funds from the Deutsche Forschungsgemeinschaft (German Research Foundation; Grant INK 12/A1, Project C). The authors gratefully acknowledge the generous assistance provided by several Western Australian bush fire fighting agencies, including the W.A. Fire Brigades, the Department of Conservation and Land Management, and the Bush Fires Board. We also thank Brent Parkin for contributing to the protocol analysis of Experiment 1, Mike Mundy for assistance with manuscript preparation, and Donald Rowe and Mike Kalish for comments on an earlier version. Special thanks are due Sarah Ryan and LihYi Soong for their assistance with data collection and analysis for Experiment 2. Correspondence concerning this article should be addressed to S. Lewandowsky, Department of Psychology, University of Western Australia, Nedlands, WA 6907, Australia (e-mail: lewan@psy.uwa.edu.au).

determining the opening of each play. Inflexibility can also be observed without artificial change to a domain, as has been shown by Marchant, Robinson, Anderson, and Schadewald (1991). Expert tax professionals and novice tax students were compared in their ability to transfer knowledge from one hypothetical tax case (called the source) to another (the target). Both cases represented exceptions to a general tax principle. Not unexpectedly, experts performed better overall than novices. However, for the experts, the extent of transfer was *reduced* when processing of the source was *enhanced* (e.g., by generating hypotheses about the relevant court ruling). Marchant et al. concluded that enhanced processing of the source primed a highly proceduralized rule in the experts' memory, which subsequently interfered with recognition of the exceptional status of the target.

Expediency, by contrast, refers to the efficiency of skill acquisition. Hecht and Proffitt (1995) identified expediency among expert waitresses and bartenders, who were asked to indicate the expected surface orientation of liquid in a tilted container. Despite their extensive relevant experience, the experts predicted surface orientations to deviate more from the horizontal than did unpracticed control subjects. This bias probably reflected an occupational emphasis on preventing spillage, which, in turn, may have focused attention during skill acquisition on the relation between liquid and container, rather than on that between liquid and environment (Hecht & Proffitt, 1995).

Overall, the analysis of performance errors and limitations contributes fundamentally to our understanding of expertise (e.g., Johnson, Grazioli, Jamal, & Zualkernan, 1992). The observed errors and limitations contravene the conventional emphasis on the apparent rich integration and flexibility of expertise (e.g., Boshuizen & Schmidt, 1992) and suggest, instead, that expertise may be divided into separate, not necessarily fully integrated components. In the case of inflexibility, experts clearly struggle to integrate new knowledge into their existing expertise. In the case of expediency, experts arguably retain some general-world knowledge (i.e., liquid surfaces are horizontal), notwithstanding their use of a different domain-specific rule (i.e., surfaces tilt with the container) in the experiment.

BUSH FIRE PREDICTION

Turning to the domain of this article, bush fire control is a candidate for expediency, because it requires the cognitive integration of several variables to predict a dynamically unfolding event. The spread of bush fires is determined by many variables—among them, ambient temperature, moisture content of the fuel, wind speed and direction, and topography (Albini, 1984). In any particular instance, assuming uniform fuel, the last two variables assume critical importance, with the wind frequently driving a fire in its direction and the slope of the terrain either accelerating (congruent slopes) or slowing

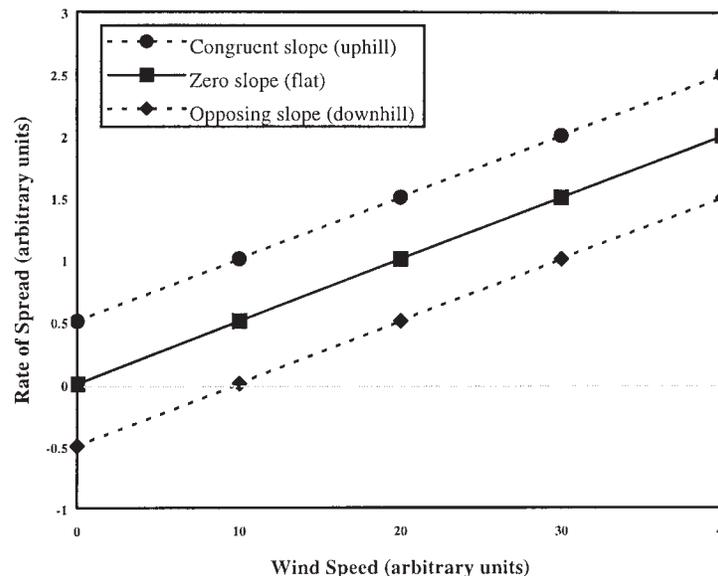


Figure 1. Rate of spread of bush fires as a function of wind speed and slope, derived from a representative mathematical model. Slope is defined relative to the direction of the wind; a congruent slope indicates that the wind is blowing uphill, and an opposing slope that the wind is blowing downhill. (The linearity between rate of spread and wind speed is an idealization that corresponds most closely to high fuel moisture levels; with dry fuels, rate of spread increases as an exponential function of wind speed.)

(opposing slopes) the spread. Figure 1 illustrates this quantitative relationship, using a representative mathematical fire model (Beck, 1995).

The figure shows the effect of wind speed on the rate of spread downwind for three different surfaces; a congruent slope (defined to be uphill with respect to wind direction), a flat terrain, and an opposing slope (downhill with respect to wind direction). For congruent slopes, wind speed and steepness have an additive effect on rate of spread. The situation is psychologically more interesting for opposing slopes. As is shown in the figure, given light winds on downhill slopes, the fire may spread *up* the slope *into* the wind, represented here by negative rates of spread.

This physical model must be known and used by bush fire commanders, whose ability to control a fire largely depends on correctly predicting its spread. How accurate, then, are those expert predictions? What is the experts' knowledge of fire, and how does it correspond to the physical model? How do experts cognitively combine the effects of wind and slope? The first experiment provided a broad investigation of these issues.

EXPERIMENT 1

Method

Overview and Design. Expert bush fire commanders participated in five simulation trials. On each trial, a fire detection point was shown on a map, and subjects were asked to think aloud while predicting the final extent of the fire and deciding how it should be brought under control. Stimulus fires were selected from archival records, and information about actual weather conditions, resource availability, and the like was available upon request.

A different and unique set of environmental conditions was associated with each fire. Analysis focused on prediction accuracy and verbal protocols.

Subjects. A sample of 14 experts was recruited from several fire fighting organizations in Western Australia: the Bush Fires Board (BFB), the Department of Conservation and Land Management (CALM), the Western Australian Fire Brigades (WAFB), and the Shire administrations of Mundaring and Swan.

The subjects had an average of 18.31 years experience (range, 3–30) and had fought an average of 27.43 bush fires during the preceding 12 months (range, 0–200). All the subjects had routinely served as fire ground commanders or in other leading decision-making roles. The selection criteria for our experts, therefore, mirrored standard practice in the local fire fighting community. Participation was voluntary and without remuneration.

Stimulus fires. Fires were selected subject to the constraints that all relevant data (e.g., weather conditions, resources and personnel deployed, true final extent of the fire) were available and that stimuli included a wide range of topographies and environmental variables. Fires, on average, had occurred 4.7 years previously (range, 3.6–5.6). Four of the fires (named F6, F7, F11, and F12) conformed to the standard pattern, in which wind determined the approximate direction of the fire. Importantly, one of the fires (F10) was anomalous, in that it spread uphill against the prevailing (light) winds.

Experimental maps. To maximize generalizability, three different types of maps were used. Two of the maps used contour lines to represent elevations and were plotted at a scale of 1:50,000. One of the contour maps, in addition to major roads and tracks, also in-

cluded hydrographic data, extensive labeling, and a variety of cadastral (i.e., land tenure and ownership) information. The third map provided a three-dimensional surface model as seen from an elevated southerly point. All the maps were approximately 1.5×1.0 m in size. Across subjects, fires were assigned to maps approximately equally often.

Procedure. The subjects were tested while on their regular duty (e.g., at fire stations). The subjects were told that they were to assume the role of fire ground commander and to think aloud while predicting how the fire was likely to spread and how it should be fought. The subjects were instructed to draw an outline of the final predicted extent.

Only the detection point was provided at the outset of each trial, and the subjects were told that any additional information (e.g., about weather and available resources) was available on explicit request. Two subjects insisted on the use of tabular material to assist in computation of quantitative fire predictions.

Trials were subject paced. All utterances were tape recorded for subsequent protocol analysis. Experimental sessions lasted approximately 1 h.

Results and Discussion

Predicted final extents. Two observations from 1 subject were discarded because wind direction was accidentally reversed (i.e., *south easterly* was interpreted as wind *from* the northwest), and another two observations from a different subject were discarded because no fire perimeter was provided. Prediction accuracy was captured by two measures: overlap between the predicted and the actual extent, and the angle between the predicted and the actual major axis.

Two measures of overlap were computed, one that took into account the total size of both the actual and the predicted extents and one that took into account only the predicted size. Because both yielded the same qualitative results, the simpler one is reported here, consisting of the ratio of the correctly predicted burnt area to the entire predicted area. Thus, a predicted area entirely contained within the actual fire would yield a score of 1, whereas if half of the predicted area overlapped with the true fire, the overlap score would be .5, and so on.

A one-way within-subjects analysis of variance (ANOVA) compared the five stimulus fires, using a single observation from each subject-trial. The analysis revealed a large main effect of fire [$F(4,48) = 9.34$, $MS_e = .06$, $p < .0001$] that was due to the anomalous fire's (F10) leading to dramatically worse performance than any of the others. The size of this effect can be visually assessed in Figure 2, which shows predictions for the anomalous fire (F10) and for a representative fire from the other set (F11).

The second prediction measure consisted of the angle between the predicted and the actual major axis, defined as the line connecting the detection point to the most distant point on the fire perimeter. Perfect predictions would correspond to an angle near zero, whereas an angle near 180° would reflect predictions maximally opposite to the actual fire spread. The analysis revealed significant differences between fires [$F(4,48) = 13.48$, $MS_e = 1253.98$, $p < .0001$]. Exploration by Newman-Keuls identified

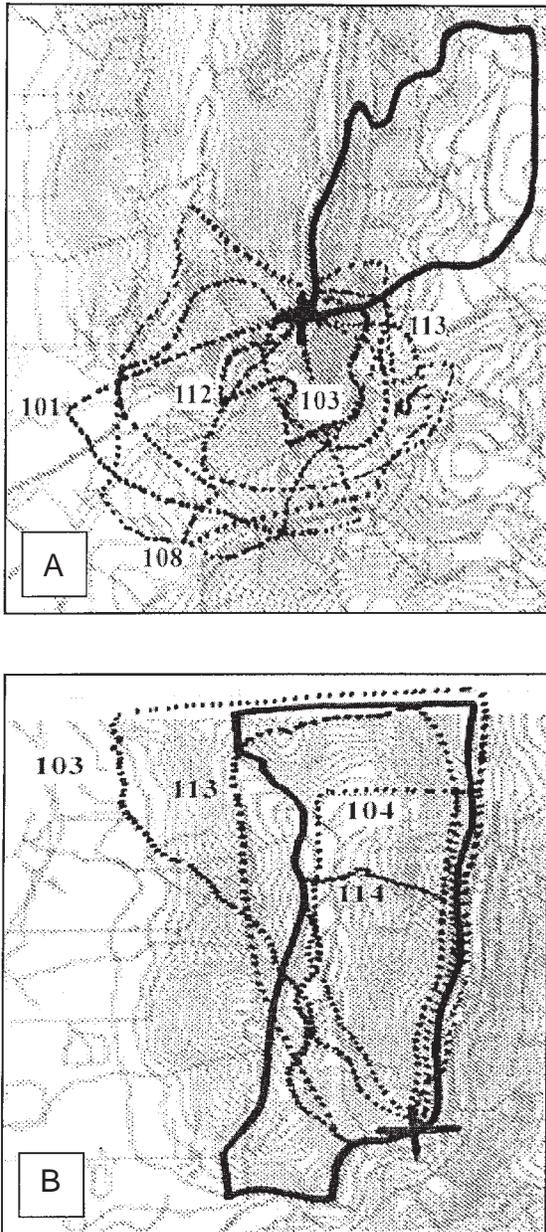


Figure 2. Prediction performance in Experiment 1 for the anomalous fire (F10, panel A) and for a typical fire (F11, panel B). Detection points are indicated by the bold cross, the actual final extents are enclosed by a solid border, and predicted final extents are shown by dashed outlines identified by participant number.

two fires (F6 and F10) that jointly differed from the others. Table 1 shows the means of both measures.

Overall, the subjects clearly struggled with the anomalous fire (F10): Prediction overlap was minimal (.09), and predicted direction of spread was very different from that actually observed (major axis angle, 97.50°). This compares with a mean for those two measures of .49 and

29.13°, respectively, for the other four fires. Protocol analysis revealed that the experts had difficulty recognizing the predominant role of slope for F10.

Verbal protocols. Protocols could be consistently divided into three consecutive segments, with little overlap: The subjects typically began by requesting information about the fire (not analyzed here); this was followed by their stating fire predictions and, finally, deciding how the fire was to be controlled. Separate analyses were conducted for utterances relating to fire prediction and fire control, with an emphasis here on the former.

Utterances relating to fire prediction were first classified into five broad categories. The subjects mentioned the effects of wind (147 utterances) and terrain (151) approximately equally often, followed by specific comments about the extent of the fire (92) and statements about fire breaks (25). An additional 25 statements referred to miscellaneous other variables (e.g., temperature, humidity). The following two statements illustrate references to the effects of wind and terrain, respectively: “with north–northwest [wind] it would burn down the gully” (Fire F6, Subject P112); “would burn up the back of hills” (F6, P112).

Further analysis was limited to the primary variables, wind and terrain. Table 2 shows the mean frequency across subjects with which each variable was mentioned. The table also shows how many subjects (out of 14) mentioned each variable first in the prediction segment of the protocols. Clearly, although wind and terrain were referred to with equal frequency overall, predictions mentioned wind *before* terrain.

To illustrate, consider first a sequence of utterances (Fire F10, Subject P101), coded by category and order (in parentheses) that relied almost entirely on wind: (*Wind 1*) “wind will take it”; (*Wind 2*) “fire goes with the wind”; (*Wind 3*) “ENE wind, means fire coming SW”; (*Specific extent 4*) “fire creeping 1/2 km SW”; (*Wind 5*) “wind swung around at 1300 to ENE”; (*Wind 6*) “wind will fan fire S until 1500”; (*Wind 7*) “wind changed to NW”; and (*Fire break 8*) “fire stopped before 2 km, provided a road or fire break there.”

Now consider a more balanced sequence for the same fire (P102): (*Wind 1*) “NE swinging NW, so fire started going SW, then swung around”; (*Specific extent 2*) “1 ton/hectare [fuel loading] and low wind speeds, fire not really doing much”; (*Terrain 3*) “1 km/hr up slope”; (*Specific 4*) “fire tends to go where fuel is available, which is behind it”; (*Terrain 5*) “fire would go up the ridges”;

Table 1
Two Measures of Prediction Performance
for All Fires in Experiment 1

Fire	Overlap	Major Axis
F6	.35	61.31
F7	.52	25.15
F10	.09	97.50
F11	.48	16.46
F12	.61	13.62

Note—F10, the anomalous fire, is show in bold.

Table 2
References to Primary Variables, Wind and Terrain,
in Verbal Protocols Across All Fires in Experiment 1

Fire	Variable			
	Wind		Terrain	
	<i>M</i>	%	<i>M</i>	%
F6	2.7	86	1.8	0
F7	2.3	79	1.6	7
F10	2.3	50	2.6	21
F11	1.6	64	2.1	21
F12	1.8	57	2.9	21
Average	2.14	67	2.17	14

Note—*M*, mean number of utterances; %, percentage of subjects who mentioned variable first. F10, the anomalous fire, is shown in bold.

(*Terrain 6*) “fire will zoom up hills either side”; (*Wind 7*) “wind swinging around would push fire down the valley”; and (*Specific 8*) “might get to the very top here.” Note that, despite repeated mention of the terrain, even in this case, wind direction explicitly determined fire direction.¹

Notwithstanding the obvious prominence of wind, the protocols also confirmed that experts were aware of the role of slope. Of the 151 total utterances relating to terrain, 42 statements implied acceleration or deceleration owing to slopes (e.g., “fire will *zoom* up hills either side” [P102, F10]; or “fire is going to *race* up side of hill” [P104, F12]), and another 24 explicitly pointed to the modulating influences of terrain (e.g., “speed would halve when at top of slope” [P113, F10]; or “come up there and then slow down” [P106, F10]). The remaining statements were neutral with respect to the effects of terrain (e.g., “fire will head toward coastal plains” [P103, F10]). Overall, every subject referred to the impact of slope in one or more protocols.

The protocols suggest that experts, by default, first derive fire direction from wind direction and then adjust the rate (but not the direction) of spread by considering slope. This *wind dominance* strategy is often successful, but, as is shown by F10, it fails systematically when wind and slope are in opposition. We now turn to analysis of the fire control segments of the protocols, which confirmed that this error did not reflect insufficient knowledge of the consequences of wind–slope opposition. Rather, the wind dominance error seemed to arise from a failure to access domain-relevant knowledge that was demonstrably available in other situations.

Fire control utterances were first divided into eight categories, consisting of terrain (21 statements), general comments about the weather (8), use of back burns (62), personnel (12), equipment (73), fire breaks (88), strategy (143), and a miscellaneous category (9). The focus here is on utterances relating to back burns, because they further illuminate the experts’ understanding of fire behavior.

Back burns are controlled burns, ignited ahead of the to-be-controlled fire, with the intention of starving the to-be-controlled fire of fuel. An obvious risk with back

burns is that they may end up traveling in the same direction as the to-be-controlled fire, thus extending rather than controlling its spread. At least three strategies exist to minimize that risk, all of which were represented among the 62 relevant statements: First, back burns can be ignited off firebreaks (e.g., on the upwind side of roads; 21 mentions). Alternatively, back burns may be lit in anticipation of a wind change (e.g., an approaching sea breeze; 2 mentions). Finally, and most relevant here, back burns can exploit the terrain: When lit on leeward slopes if winds are light, back burns will spread uphill and into the wind to meet the to-be-controlled fire approaching on the windward side. This technique, physically isomorphic to the situation involving F10, was mentioned four times in the protocols, by Subjects P104 and P109, confirming that at least some experts were aware that slope may override opposing wind.

Intriguingly, these two individuals nonetheless mispredicted the direction of F10 by 50° (P104) and 135° (P109), suggesting that their predictions were a function not only of physical variables, but also of the context and the history of the particular fire. One subject revealed this context specificity in a particularly striking utterance: “Having this at the bottom and burning into wind [the back burn], it’ll be a nice slow burn up, and this [the to-be-controlled fire] will be a nice slow burn down, because it’s burning down the hill with the wind behind it” (P109, F10). The utterance is indicative of extensive prior experience with wind-driven to-be-controlled fires and with back burns that spread uphill, but only one of the two opposing predictions can be applicable to the prevailing physical conditions.

It is possible that other subjects might have shown similar contradictory predictions, if the opportunity had arisen. That is, other subjects who failed to predict F10 might have made the correct predictions for back burns, but this was not discovered, because those subjects either chose not to back burn or, if they did, chose not to exploit slope. This possibility was supported by a follow-up interview conducted a considerable time after the study. During the interview, 3 out of 5 of the subjects who could be contacted spontaneously cited the effect of slope on back burns.

Summary and limitations. This experiment yielded two principal findings. First, in relying on the wind dominance strategy identified by the protocols, experts provided several accurate predictions alongside a strikingly incorrect one. The erroneous prediction occurred when the two primary predictor variables, wind and slope, were in opposition.

The second principal finding, perhaps more intriguing but also more tentative, concerned the context specificity of expert performance. Two subjects who exhibited erroneous wind dominance when predicting F10 nonetheless predicted back burns to spread uphill and into the wind. In one instance, an expert predicted the

physically impossible situation of two fires moving toward each other on the same slope, each driven by the predictor variable most relevant to the context of the fire.

There are limitations associated with both findings: First, the anomalous fire differed from the remaining stimuli in several ways other than the opposition of wind and slope (e.g., location, temperature). The wind-dominance error may, therefore, reflect one of those uncontrolled incidental differences. In support of this possibility, another fire, F6, was statistically identical to F10 for the major axis measure. Second, context specificity was only observed in 2 subjects and without knowledge of the physical circumstances surrounding their predicted back burns. Both of those limitations were addressed in Experiment 2.

EXPERIMENT 2

Stimuli for Experiment 2 were arbitrarily constructed to obtain experimental control over slope, wind, and the orientation of those two predictor variables. On the basis of the first experiment, we expected experts to use a wind dominance strategy throughout, with prediction errors ensuing whenever wind opposed slope.

Likewise, Experiment 2 varied the context (back burning or fire fighting) in which an otherwise identical fire was presented. If context specificity were a general characteristic of expertise in this domain, predictions might differ between contexts for fires that were otherwise physically identical.

An additional purpose of Experiment 2 was to ascertain how the predictions of experts differed from those of novices. In particular, any claim that the present findings revealed a characteristic attribute of expert knowledge must be accompanied by a demonstration that the experts differed from novices in identifiable ways.

Method

Overview and Design. The experiment compared three groups of subjects of varying domain-relevant expertise: The novice group was totally ignorant of bush fires and lacked practice with maps; the orienteering group included skilled map readers without any fire fighting expertise; and the expert group consisted of experienced bush fire fighters. This between-subjects variable was fully crossed with two orthogonal within-subjects manipulations, consisting of wind-slope orientation (opposing or congruent) and fire context (back burning or fire fighting).

Stimulus fires were constructed by arbitrarily placing detection points at the bottom of steep slopes and arbitrarily assigning weather conditions. Across subjects, all the fires appeared equally often in all within-subjects conditions. On each trial, the detection point was presented on a map, and the subjects were asked to think aloud while predicting the spread of the fire. Unlike the first experiment, all the relevant weather information was available throughout. Analysis was limited to performance variables.

Subjects. A total of 24 subjects participated, with 8 recruited for each of the three groups. All the subjects participated voluntarily.

The novice group consisted of first-year psychology students at the University of Western Australia, who were screened to be ignorant of bush fires. The subjects in this group participated in exchange for credit toward a course requirement.

The subjects in the orienteering group were members of the Orienteering Association of Western Australia, with an average of

15 years orienteering experience (range, 12–25). Orienteering is a sport in which competitors use a map and compass to find their way through unfamiliar bush territory. The subjects in this group, therefore, were expert at reading and interpreting maps but had no direct knowledge of bush fires. The subjects in this group received \$5 (Australian) in exchange for participating.

The expert group consisted of members of the Western Australian Fire Brigades, who participated without remuneration. The experts had an average of 11.4 years bush fire fighting experience (range, 3–24 years) and had fought an average of 16.9 bush fires during the preceding 12 months (range, 2–50). All the experts had been involved in bush fire decision-making roles during their career.

Stimulus fires. Stimuli were constructed by combining eight detection points with arbitrary weather information. One of the detection points was identical to the anomalous fire (F10) of Experiment 1; the remainder were chosen to be at the bottom of steep slopes spread across a relatively large (200 km²) region. With the exception of the point replicating F10, the fires were placed away from major natural fire breaks, such as roads, rivers, or railway tracks. This ensured that the fire, in reality, could have spread in any direction from the detection point.

Each fire was accompanied by a date (always during the fire season), detection time (typically around 11 a.m.), temperature (mean, 28.3°C; range, 22°–35°), humidity (mean, 33.5%; range, 22%–58%), and wind speed. Winds were light (mean, 3.9 km/h; range, 1–8) to ensure that slope would be the dominant variable if it opposed wind in reality. To maximize realism, some fires were accompanied by changing weather conditions (e.g., a minor shift in wind speed and direction within an hour of fire detection).

Regardless of the condition in which a fire appeared, weather information remained unchanged, except that wind either opposed slope (with a deviation from exact opposition of $\pm 10^\circ$ to maximize realism) or was congruent with slope (with the same range of deviation). Regardless of wind-slope orientation, when a stimulus fire appeared in the back burn context, the to-be-controlled fire was represented by a filled irregular shape of approximately 2–6 km² size located at the top of the slope above the point of origin of the back burn. To avoid demand characteristics, the shape of the to-be-controlled fire did not suggest a primary direction of spread.

Fires were presented on segments of the detailed 1:50,000 contour map used in Experiment 1. Map segments were color copies (size, 210 × 297 mm) of the original. Wind direction was represented by an arrow and alphabetically, using a 16-point compass. Each map was accompanied by the appropriate scale.

The eight stimulus fires were randomly split into two sets of four, with each set being presented to a random half of the subjects in each group. Across subjects within a given group, each fire was presented once in each condition. For all the subjects, the trials in the fire fighting context preceded back burns, and wind opposed slope for the first fire within each context.

Procedure. The experts participated while on duty. All the other subjects were tested in a university laboratory. All the subjects were tested individually. The subjects were told to predict the spread of various fires by drawing an outline of their predicted final extent. In the fire fighting context, the subjects were to assume that 2 h had elapsed before fire control measures were initiated. For back burns, the subjects were told that the back burn was intended to fight the to-be-controlled fire. If the subjects queried fire control, they were informed that the back burn was a controlled fire (the experts automatically assumed back burns to be controlled).

The subjects were not informed about the exact purpose of the experiment. The four trials were subject paced, and experimental sessions lasted approximately 15–20 min.

Results and Discussion

Because actual outcomes did not exist, analysis focused on predicted fire spread relative to wind direction.

Each predicted extent was numerically summarized by measuring the distances (in centimeters) from the point of detection to the downwind and upwind boundaries. If wind direction changed during the prediction interval, the average of the two directions was used as the reference axis. Two analyses were conducted. The first one considered the sum of both distances (corresponding to the total spread along the wind axis). This served as an index of the overall predicted size of a fire. The second analysis considered the ratio of the spread downwind to the total distance along the wind axis. This *downwind ratio* was insensitive to fire size but revealed the predicted direction of fire spread.²

Total distance. Total predicted distances along the wind axis were entered into a 3 (group) × 2 (fire context) × 2 (wind-slope orientation) between-within-subjects ANOVA. The analysis revealed significant main effects of group [$F(2,21) = 11.97, MS_e = 6.57, p < .01$] and fire context [$F(1,21) = 66.81, MS_e = 6.00, p < .0001$], as well as a significant interaction between those two variables [$F(2,21) = 4.75, MS_e = 6.00, p < .05$]. Wind-slope orientation did not have an effect and did not contribute to any interactions (all $F_s < 1$). The means underlying the significant effects are shown in Table 3.

Predicted sizes clearly decreased with expertise: The differences were most extreme in the fire fighting context, where the novices predicted a wind axis length of 8.62 cm (corresponding to 4.31 km in reality) and the experts predicted 3.63 cm (1.82 km in reality). Despite the absence of objective performance measures, the plausibility of the predictions can be assessed by comparison with the known outcomes in Experiment 1. The longest wind axis among those actual fires was around 2.5 km (F11), obtained with a wind speed of 10 km/h, more than double the average wind speed (3.9 km/h) in Experiment 2. It follows that the novices widely overestimated the extent of fires (nearly double the length of F11 with half the wind speed), whereas the experts' predictions appeared well calibrated (their predicted 1.82 km with 3.9 km/h winds is compatible with 2.5 km and 10 km/h for F11).

The table also shows that the predicted sizes of to-be-controlled fires were greater than the predicted extents of back burns, with the magnitude of that difference again decreasing with increasing expertise. In part, the smaller back burns may have reflected subjects' knowledge that these were controlled fires.

Table 3
Total Predicted Length (in Centimeters) of Fires Along Wind Axis in Experiment 2

Group	Context		Average
	Fire Fighting	Back Burning	
Novices	8.62	2.54	5.58
Orienteers	5.64	1.79	3.71
Experts	3.63	1.30	2.47
Average	6.00	1.88	

Downwind ratio. Downwind ratios greater than .5 represented fires that spread primarily with the wind, whereas ratios of less than .5 referred to fires that spread primarily into the wind. Downwind ratios were set to .5 (indeterminate) for four cases (all involving experts during back burning) in which the predicted outline consisted of a single line drawn orthogonally to the wind axis. The conclusions remain unchanged if those observations are, instead, assumed to be missing.

An ANOVA on downwind ratios again showed context to have an effect [$F(1,21) = 18.86, MS_e = 0.03, p < .001$], with back burns spreading downwind to a lesser extent (.70) than to-be-controlled fires (.86). There was also a main effect of wind-slope orientation [$F(1,21) = 8.41, MS_e = 0.03, p < .01$], with congruent fires spreading downwind more (.83) than fires in which slope opposed wind (.73).

All three pairwise interactions were significant [context × group, $F(2,21) = 5.87, MS_e = 0.03, p < .01$; orientation × group, $F(2,21) = 8.34, MS_e = 0.03, p < .01$; context × orientation, $F(1,21) = 7.43, MS_e = 0.03, p < .015$] but were qualified by a significant interaction of all three experimental variables [$F(2,21) = 4.48, MS_e = 0.03, p < .03$]. Table 4 shows the underlying means, which provide the focus for discussion.

The table suggests that the three-way interaction and, by implication, the subordinate two-way interactions arose from the large difference between congruent and opposing wind-slope orientations for experts in the back burning context (boldfaced entries in Table 4). That is, whereas the two nonexpert groups ignored slope regardless of context, the experts ignored that predictor (and exhibited a wind-dominance error) only in the fire fighting context. When back burning, the experts very plausibly predicted the fire to spread uphill and into the wind.

This contention is supported statistically by three within-subjects ANOVAs conducted for each group separately, using wind-slope orientation and context as the only variables. Only the analysis for experts yielded significant results, with a main effect of context [$F(1,7) = 28.29, MS_e = 0.31, p < .0001$], wind-slope orientation [$F(1,7) = 18.20, MS_e = 0.041, p < .0001$], and an orientation × context interaction [$F(1,7) = 10.65, MS_e = 0.047, p < .02$]. Importantly, that interaction failed to reach significance for the novices [$F(1,7) = 5.34, MS_e = 0.010, .05 < p < .10$] and the orienteers ($F < 1$). The nearly significant effect for the novices reflected the paradoxically *greater* downwind spread of to-be-controlled fires when wind opposed slope and will not be considered further. For the two nonexpert groups, all main effect F_s were ≤ 1.60 .

An examination of individual differences among the experts revealed that, in the back burning context, only 1 subject predicted the fire to spread downwind, with a downwind ratio of .67. This contrasted sharply with the fire fighting context, in which the minimum individual downwind ratio was also .67, with all but 1 subject yield-

Table 4
Downwind Ratio for Predicted Fire Spreads in Experiment 2

Group	Wind-Slope Orientation	Context		Average
		Fire Fighting	Back Burning	
Novices	congruent	.79	.80	.79
	opposing	.87	.71	.79
Orienteers	congruent	.84	.75	.79
	opposing	.82	.78	.80
Experts	congruent	.95	.87	.91
	opposing	.90	.32	.61

Note—Values above .5 represent fires that spread primarily with the wind, whereas values below .5 represent fires that spread into the wind.

ing greater ratios. Hence, only 2 of the 16 individual downwind ratios for experts were identical between contexts, notwithstanding the physical identity of the fires.

The individual difference analysis also implies that, in contrast to Experiment 1, all the experts exhibited the wind dominance error in the fire fighting context. This raises to 20 the total number of experts (out of 22 across both studies) who showed that error.

Summary. The results of Experiment 2 are readily summarized. First, we replicated the wind dominance error observed with the anomalous fire (F10) in Experiment 1. Because all the stimuli served in all the conditions, this clearly identified wind-slope orientation as the single critical variable. It therefore seems to be a general attribute of expertise in bush fire fighting that slope is ignored when determining fire direction, even when, in reality, slope is likely to be the more important predictor. Few experts (about 1 in 10 in our samples) appear able to discount wind and to base their predictions on slope alone.

At the same time, there were two clear indications that the experts were, in fact, expert and superior in their predictions to the other two groups. First, when predicting the size of fires, experts were better at calibration than the orienteers or novices. Second, whereas novices and orienteers consistently ignored slope, the experts correctly exploited slope, where appropriate, when predicting the spread of back burns.

The latter finding is most intriguing, because it confirms the suggestion of Experiment 1 that expert performance is context specific. A remarkable consequence of that context specificity was that, when an identical fire was presented in two different contexts on different occasions, experts made exactly opposing and mutually exclusive predictions.

GENERAL DISCUSSION

We first briefly discuss the wind dominance error and differentiate it from other expert limitations. We then focus on the context specificity of the wind dominance strategy and show that it differs from previous findings involving context and expertise.

The Wind Dominance Error

The experts performed poorly in the present experiments in circumstances that were domain relevant and

not totally uncommon. There is every reason to assume that the experts had encountered a fire in which wind opposed slope before participating in the study. However, that said, given the topography and prevailing winds of coastal Western Australia, wind and slope are more likely to be congruent than opposing. One senior expert in Experiment 1 (P116, who provided correct predictions for F10) estimated that wind and slope were congruent for 90% of the fires in the region. Thus, an explanation for our results should first consider expediency, similar to that observed with waitresses and bartenders by Hecht and Proffitt (1995).

On this view, experts acquired a simple rule, the wind dominance strategy revealed through protocols in Experiment 1, in preference to mastering the physically correct but more complex alternative. The expedient wind dominance strategy provided good predictions for most (congruent) fires but failed when wind and slope were in opposition. This view is compatible with the lack of transfer between exceptional legal cases observed by Marchant et al. (1991), which demonstrated the supremacy of strong rules over exceptions, even if exceptions were given prominence in the experiment.

However, unlike the studies by Hecht and Proffitt (1995) and Marchant et al. (1991), the expedient wind dominance strategy was not uniquely associated with expertise. In Experiment 2, all the subjects failed to discount wind when it was in opposition to slope. Considered in isolation, the wind dominance error, therefore, did not illuminate the unique knowledge of experts and how they differ from novices. Rather, experts were qualitatively differentiated from novices by the context specificity of their performance.

Context Specificity of Expert Performance

Back burns obey the same physical rules as any other fire. Accordingly, the two nonexpert groups in Experiment 2 did not distinguish between back burns and to-be-controlled fires and provided qualitatively identical (wind-dominant) predictions in both contexts. The experts, by contrast, made qualitatively different—and contradictory—predictions. Clearly, therefore, context-specificity was selectively associated with expertise. This finding differs from previous reports of context effects.

Context and domain specificity. At the outset, we cited specificity of skill as a defining attribute of expertise. For example, training can selectively raise digit span to 80 without also increasing the capacity to encode letters and other materials (Chase & Ericsson, 1981). Likewise, the chess master's superior memory for naturally occurring chess configurations is dramatically reduced, to levels barely above that of novices, when board positions are randomized (Chase & Simon, 1973; Gobet & Simon, 1996).

Overall, there are numerous reports that *context*, used here to refer to the characteristics of the problem and its materials, is a major determinant of expert performance: Whenever context is changed, so that a problem is taken outside the expert's domain, however narrowly defined,

performance greatly deteriorates. Critically, that type of context effect has little or no bearing on the present results: The two experiments in this article manipulated context *within* a domain, by alternating between-problem situations with which the experts were thoroughly familiar.

Context shifts within a domain. Little is known about the effects of a context shift within a domain of expertise. One relevant line of research has focused on people who lack formal schooling but are able to solve mathematical problems in everyday contexts (see, e.g., Carraher, Carraher, & Schliemann, 1985; Nunes, Schliemann, & Carraher, 1993). The subjects in these investigations of “street mathematics” included young Brazilian street vendors, fishermen, construction foremen, and cooks, all of whom had minimal formal schooling but were highly competent at solving mathematical problems associated with their domain of expertise.

Of greatest interest here is a context manipulation involving expert cooks (Schliemann & Carraher, 1993). Subjects were presented with identical proportionality problems, either in a pricing context (“If 2 kg of rice cost 5 cruzeiros, how much do you have to pay for 3 kg?”) or in a recipe context (“To make a cake with 2 cups of flour you need 5 spoonfuls of water; how many spoonfuls do you need for 3 cups of flour?”). Importantly, both problem contexts were familiar to the subjects and relevant to their domain of expertise. Schliemann and Carraher reasoned that social convention dictated accuracy in the pricing context, whereas estimation might be acceptable for recipes.

Those expectations were confirmed. In the pricing context, subjects used a variety of identifiable mathematical strategies rather than estimation, with the result that accuracy was in excess of 90%. In the recipe context, by contrast, accuracy was dramatically lower (20%), and half of the responses were based on estimation. This performance pattern—in particular, the use of different ways to solve a problem in different contexts—is analogous to the context specificity observed with bush fire experts. One potentially important qualification is that the cooks examined by Schliemann and Carraher (1993), unlike bush fire experts, had little or no formal training in the type of problems on which they were tested. Indeed, it could be argued that mental arithmetic, although relevant to the cooks’ domain of expertise, did not itself form part of the domain.

Context specificity in bush fire prediction. The present findings, then, are unique in two important ways. First, both domain-relevant contexts were part of the training repertoire of bush fire experts, and both regularly occur in the field. Second, the context shift resulted not only in a reduction of performance, so that experts performed worse in one context than in another, but in a qualitative reversal of predictions, so that the same physical situation gave rise to two mutually exclusive and contradictory predictions.

Thus, the principal empirical contribution of this article is that context plays a major role not only in the retrieval of explicit (see, e.g., Godden & Baddeley, 1975) and implicit (see, e.g., Bainbridge, Lewandowsky, & Kirsner, 1993) information, but also in the application of expert knowledge.

Implications for Views of Expert Knowledge

We explore two possible explanations for the observed context specificity. The first one extends the earlier expediency account. The second one, which we ultimately favor, views expertise as partitioned into independent parcels that are accessed depending on context.

Expediency. To extend the earlier expediency account, a simple wind dominance strategy must be augmented by a provision for context. This provision turns out to defy expedient summary, because it must explain not only the contradictory predictions when wind and slope were in opposition, but also the fact that predictions were identical across contexts when wind and slope were congruent. Thus, a plausible minimum set of expedient rules would be: Wind determines fire direction, unless the fire is a back burn and wind opposes slope, in which case the fire spreads uphill and into the wind.

There are at least two difficulties with this account: First, the expedient rules are no longer much simpler than the physical process they supposedly supplant. Second, these expedient rules do not easily capture another quantitative aspect of the results—namely, that back burns, regardless of their predicted direction, were uniformly expected to be smaller than to-be-controlled fires (Experiment 2). We, therefore, favor another explanation, based on the idea that expertise may not be fully integrated, an idea we refer to as *knowledge partitioning*.

Knowledge partitioning. The strong assumption that expert knowledge is highly integrated echoes throughout the literature. For example, according to Bédard and Chi (1992), “experts have more and stronger links among concepts, suggesting that there is a greater degree of connectedness and cross-referencing, and the pattern of connections and cross-referencing can result in a better structure” (p. 136). Glaser (1996) pointed to the centrality of the “acquisition of well-organized and integrated knowledge For the development of expertise, knowledge must be acquired in such a way that it is highly connected and articulated, so that inference and reasoning are enabled” (p. 306). Finally, it has been assumed that, during training, experts notice inconsistencies in their current knowledge, “which in turn will serve as a stimulus for further analysis . . . until an acceptable reintegration . . . is attained” (Ericsson, 1996, p. 38).

Clearly, it has been a central tenet of theorizing that expertise is characterized by notable integration, cross-referencing, and consistency. In consequence, there has been little theoretical impetus to search for inconsistencies or contradictory knowledge among experts. We know

of only two relevant results: the finding by Schliemann and Carraher (1993) that cooks use different strategies—and produce different solutions—when solving identical mental arithmetic problems in different contexts, and the present results involving bush fire experts.

These inconsistencies suggest to us that expertise, whether acquired expediently or not, is tied to the context of its acquisition and utilization without necessarily being cross-referenced to preexisting knowledge. Notwithstanding the predominant theoretical emphasis on integration, we suggest that knowledge may be partitioned and that even major inconsistencies or contradictions may persist during expertise acquisition.

Our view is related to the knowledge encapsulation framework proposed by Boshuizen and Schmidt (e.g., 1992). Knowledge encapsulation is defined “as the progressive subsumption, or packaging, of lower level concepts ... under a limited number of high-level concepts with the same explanatory power” (Schmidt & Boshuizen, 1993, p. 347), which allows access of an entire cluster of related facts through a single higher level concept.

We share the view that clusters of facts can be accessed through higher level concepts (although we remain neutral as to some of the other specific attributes of encapsulation, as defined by Schmidt & Boshuizen, 1993). We additionally assume that components of knowledge are hidden from each other when encapsulated in separate *packages*. In consequence, although each package by itself may be well integrated, the invisibility of components across packages may cause contradictions and inconsistencies.

Specifically, we suggest that bush fire controllers may learn to condense all knowledge relevant to back burning in one package and all knowledge about to-be-controlled fires in another package, and so on. Assuming that each package mirrors the constraints of the training environment, the back-burning package will likely contain the fact that fire spreads uphill and into the wind (because back burns are frequently lit with that fact in mind). Conversely, the package devoted to fire prediction more than likely will contain the knowledge that wind determines fire direction and that slope will affect speed but not direction (because most to-be-controlled fires conform to that pattern).

Although each package may be highly integrated, the continued encapsulated usage of that knowledge, partitioned by context, prevents any resolution of inconsistencies or contradictions. Thus, experts may predict that two fires, placed on the same slope under identical wind conditions, will spread in opposite directions and toward each other.

Identifiability of knowledge partitioning. Knowledge partitioning currently is not a complete model of expertise. Rather, it is best understood as a novel framework, stimulated by the present results, that was proposed to counter the widespread emphasis on the seemingly unlimited integration of expertise. We believe that emphasis has detracted attention from the search for interesting inconsistencies in expert behavior.

We are not presently committed to a detailed model of the knowledge parcels; thus, they could be instantiated as single productions in a production system such as ACT (see, e.g., Anderson, 1993) or as separate modules within a gated connectionist model, such as ATRIUM (Erickson & Kruschke, 1998). Regardless of their instantiation, we believe that access to knowledge parcels is guided by problem context and that different contexts may elicit different knowledge for an identical problem.

Empirically, knowledge partitioning is most readily identified through inconsistencies or contradictions across contexts (see Sloman, 1996, for examples of the theoretical diagnosticity of contradictions). However, relying on contradiction alone is potentially circular, because the inferred explanation (i.e., separate knowledge parcels) is also the to-be-explained result (i.e., contradiction). This circularity can be avoided through the development of techniques for the independent identification of knowledge parcels.

A candidate technique can be foreshadowed by considering a recent reexamination of the Power Law of practice by Delaney, Reder, Staszewski, and Ritter (1998). Delaney et al. showed that the ubiquitous Power Law does not always describe overall task performance but, instead, needs to consider specific strategies that subjects use to perform the task. Delaney et al. discovered that the autocorrelations among residuals (which suggest the omission of relevant predictor variables) were minimized when separate power functions were fit to sets of trials on which subjects used different strategies. This suggests that learning is specific to the strategy being used and does not include the task per se. In the present context, knowledge parcels would be identified if performance across learning trials were best described by separate power functions for sets of trials (and contexts) involving different putative parcels. Critically, in that case, parcels would be identified independently of the contradictions they might engender.

Avoiding partitioned knowledge. In a training study conducted by Lewandowsky, Dunn, Kirsner, and Randell (1997), knowledge partitioning was avoided through the use of integrative instructions. In one condition, novices learned to control simulated bush fires without specific instructions. These subjects developed the wind dominance typical of the experts in the present experiments. In another condition, Lewandowsky et al. presented subjects at the outset with an integrative diagram that summarized the physics of fire (by using the angle between the flames and the ground as the common referent for both wind- and slope-driven fires). The diagram raised sensitivity to slope and dramatically increased overall accuracy. Thus, knowledge partitioning can be avoided if instructions or diagrams foster integration of predictors from the outset of training.

CONCLUSIONS

In two studies, experts had to predict a moderately complex physical process involving two primary predic-

tor variables. The data suggested that most experts did not master the interaction between the predictors but, instead, learned to associate each with a different level of a third variable, context, that was totally irrelevant to the physical process. In consequence, experts, *because* of their expertise, made completely opposing predictions for an identical physical situation, depending on the context of the problem.

We prefer an explanation that views expertise as being partitioned, whereby conflicting information might co-exist in separate parcels of knowledge. The widely touted view that expertise is highly integrated may be characteristic of each parcel but need not hold across the entire landscape of expertise.

REFERENCES

- ALBINI, F. A. (1984). Wildland fires. *American Scientist*, **72**, 590-597.
- ANDERSON, J. R. (1993). *Rules of the mind*. Hillsdale, NJ: Erlbaum.
- BAINBRIDGE, J. V., LEWANDOWSKY, S., & KIRSNER, K. (1993). Context effects in repetition priming are sense effects. *Memory & Cognition*, **21**, 619-626.
- BECK, J. A. (1995). Equations for the forest fire behaviour tables for Western Australia. *CALM Science*, **1**, 325-348.
- BÉDARD, J., & CHI, M. T. H. (1992). Expertise. *Current Directions in Psychological Science*, **1**, 135-139.
- BOSHUIZEN, H. P. A., & SCHMIDT, H. G. (1992). On the role of biomedical knowledge in clinical reasoning by experts, intermediates, and novices. *Cognitive Science*, **16**, 153-184.
- CAMERER, C. F., & JOHNSON, E. J. (1991). The process-performance paradox in expert judgment: How can experts know so much and predict so badly? In K. A. Ericsson & J. Smith (Eds.), *Towards a general theory of expertise* (pp. 195-217). Cambridge: Cambridge University Press.
- CARRAHER, T. N., CARRAHER, D. W., & SCHLIEMANN, A. D. (1985). Mathematics in the streets and in schools. *British Journal of Developmental Psychology*, **3**, 21-29.
- CHASE, W. G., & ERICSSON, K. A. (1981). Skilled memory. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 141-189). Hillsdale, NJ: Erlbaum.
- CHASE, W. G., & SIMON, H. A. (1973). Perception in chess. *Cognitive Psychology*, **4**, 55-81.
- DELANEY, P. F., REDER, L. M., STASZEWSKI, J. J., & RITTER, F. E. (1998). The strategy-specific nature of improvement: The power law applies by strategy within task. *Psychological Science*, **9**, 1-7.
- ERICKSON, M. A., & KRUSCHKE, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General*, **127**, 107-140.
- ERICSSON, K. A. (1996). The acquisition of expert performance: An introduction to some of the issues. In K. A. Ericsson (Ed.), *The road to excellence: The acquisition of expert performance in the arts and sciences, sports and games* (pp. 1-50). Hillsdale, NJ: Erlbaum.
- ERICSSON, K. A., & CHARNES, N. (1994). Expert performance: Its structure and acquisition. *American Psychologist*, **49**, 725-747.
- ERICSSON, K. A., CHASE, W. G., & FALON, S. (1980). Acquisition of a memory skill. *Science*, **208**, 1181-1182.
- ERICSSON, K. A., & LEHMANN, A. C. (1996). Expert and exceptional performance: Evidence of maximal adaptation to task constraints. *Annual Review of Psychology*, **47**, 273-305.
- FRENSCH, P. A., & STERNBERG, R. J. (1989). Expertise and intelligent thinking: When is it worse to know better? In R. J. Sternberg (Ed.), *Advances in the psychology of human intelligence* (Vol. 5, pp. 157-188). Hillsdale, NJ: Erlbaum.
- GLASER, R. (1996). Changing the agency for learning: Acquiring expert performance. In K. A. Ericsson (Ed.), *The road to excellence: The acquisition of expert performance in the arts and sciences, sports and games* (pp. 303-311). Hillsdale, NJ: Erlbaum.
- GOBET, F., & SIMON, H. A. (1996). Recall of rapidly presented random chess positions is a function of skill. *Psychonomic Bulletin & Review*, **3**, 159-163.
- GODDEN, D. R., & BADDELEY, A. D. (1975). Context-dependent memory in two natural environments: On land and under water. *British Journal of Psychology*, **66**, 325-331.
- HECHT, H., & PROFFITT, D. R. (1995). The price of expertise: Effects of experience on the water-level task. *Psychological Science*, **6**, 90-95.
- JOHNSON, P. E., GRAZIOLI, S., JAMAL, K., & ZUALKERNAN, I. A. (1992). Success and failure in expert reasoning. *Organizational Behavior & Human Decision Processes*, **53**, 173-203.
- LEWANDOWSKY, S., DUNN, J. C., KIRSNER, K., & RANDELL, M. (1997). Expertise in the management of bush fires: Training and decision support. *Australian Psychologist*, **32**, 171-177.
- MARCHANT, G., ROBINSON, J., ANDERSON, U., & SCHADEWALD, M. (1991). Analogical transfer and expertise in legal reasoning. *Organizational Behavior & Human Decision Making*, **48**, 272-290.
- NUNES, T., SCHLIEMANN, A. D., & CARRAHER, D. W. (1993). *Street mathematics and school mathematics*. Cambridge: Cambridge University Press.
- SCHLIEMANN, A. D., & CARRAHER, D. W. (1993). Proportional reasoning in and out of school. In P. Light & G. Butterworth (Eds.), *Context and cognition: Ways of learning and knowing* (pp. 47-73). Hillsdale, NJ: Erlbaum.
- SCHMIDT, H. G., & BOSHUIZEN, H. P. A. (1993). On the origin of intermediate effects in clinical case recall. *Memory & Cognition*, **21**, 338-351.
- SLOMAN, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, **119**, 3-22.
- STASZEWSKI, J. J. (1993). A theory of skilled memory. In *Proceedings of the 15th Annual Conference of the Cognitive Science Society* (pp. 971-975). Hillsdale, NJ: Erlbaum.
- STERNBERG, R. J. (1996). Costs of expertise. In K. A. Ericsson (Ed.), *The road to excellence: The acquisition of expert performance in the arts and sciences, sports and games* (pp. 347-354). Hillsdale, NJ: Erlbaum.
- STERNBERG, R. J., & FRENSCH, P. A. (1992). On being an expert: A cost-benefit analysis. In R. Hoffman (Ed.), *The psychology of expertise: Cognitive research and empirical AI* (pp. 191-203). New York: Springer-Verlag.

NOTES

1. Two subjects made qualitatively more accurate predictions for F10 than did the remaining subjects. One of those experts (P114), provided the only utterances in which the effect of wind was explicitly discounted: "Despite easterly wind, fire will go up the slope"; "probably on other [east] side of valley system, despite easterly wind"; "Easterly too low to have any impact on rate of travel of fire." That subject provided a small angle (10°) between predicted and actual major axes (viz., mean of 97.50°).

The final subject (P116) did not explicitly discount wind but, nonetheless, predicted F10 correctly, with a major axis angle of 30°. This subject suggested that the slope turned the wind direction around, thus rendering wind and slope congruent. This may occur in reality when the heat associated with large fires generates anabatic (i.e., up-slope) drafts irrespective of the prevailing overall wind direction.

2. Fire predictions were summarized in two additional ways: First, distances from the point of detection to the most distant point on the boundary on either side of the wind axis were obtained, and second, the angle between the lines that connected those points with the detection point was measured. The angle is best understood as capturing the aspect ratio of the predicted fire outline, with smaller angles representing fires that are more spread out along the wind axis. These additional measures were found to be highly correlated with total wind length and downwind ratio, respectively, and therefore are not reported separately.

(Manuscript received January 27, 1998;
revision accepted for publication February 24, 1999.)