how surprised we are by an outcome (e.g., a court decision is not surprising if we cannot imagine it having turned out differently), how we assign blame (e.g., a person is guilty if they should have behaved differently), and when we experience regret (e.g., we regret not having behaved in some other fashion).

The analysis of event mutability could be enhanced by consideration of dependency structure. The temporality effect (Miller & Gunesagaram, 1990; Byrne, Culhane, & Tasso, 1995) is the finding that people consider the second of two independent events in sequence to be more mutable. For example, we have asked people to consider the following problem:

> Imagine two individuals (Jones and Cooper) who are offered the following very attractive proposition. Each individual is given one yellow ball and one green ball, and asked to choose one of them. If the two individuals pick balls with the same color, then each individual wins $1000. However, if they pick balls of different colors, neither individual wins anything. During this game, they are not allowed to see each other's choices. Jones goes first. Jones picks a yellow ball because Jones prefers yellow. Cooper goes next and picks a green ball because Cooper prefers green. Thus, the outcome is that neither individual wins anything.
>
> What could have been different that would have allowed Jones and Cooper to win $1000? Please circle one of the following two options:
>
> Jones could have picked a green ball. Cooper could have picked a yellow ball.

Replicating the standard result, we find that most people choose Cooper, the second actor in the problem, over Jones, the first. Byrne et al. (1995) have argued that the second of two sequential but otherwise independent events is more mutable because the first serves as an initial anchor in the construction of a mental model of the sequence. Whether or not this is the case, we find it less than adequate as an explanation of the effect, because it fails to tell us why the first and not the second event serves as a conceptual anchor. We suggest that the difference is in the encoded dependency structure. The relevance of the second event depends on the first in a way in which the relevance of the first does not depend on the second. The initial interpretation of the second event is governed by the outcome of the first event, but the initial interpretation of the first event cannot depend on the second because the second is, initially, unknown. Because more depends on the first event than the second, the second is more mutable.

The current analysis of object representation shares with the analysis of counterfactual events the assumption that the world is interpreted not just in terms of how it is, but also in terms of how it could be. People perceive objects not as static entities but in terms of their dispositions to change. This is clearly true of objects which are intrinsically dynamic, like clouds, tumors, trees, and hurricanes (cf. Leyton, 1988). Consequently, the analyses of objects and events share a concern with the determinants of mutability. However, they also display a critical difference. Kahneman and Varey (1990) champion a distinction between dispositions, causes which serve as the background for an event, and propensities, causes which determine developments during the event itself. The bases of counterfactuals are propensities; a counterfactual outcome is one that was almost produced by the causal sequence intrinsic to an event. However, the concern of this paper is the dispositions of objects, the kinds of mutations that background knowledge supports, not with the propen-