ables that can help to predict feature relevance in conceptual combination. Another critical constraint is that the dependency structure of the head must be able to support the mapped feature. The noun-noun combination *frog car* is more likely to be interpreted as a green car than as a car that hops because the feature *can hop* depends on other features that cars do not have (Love, 1996). Hence, knowing the mutability of head features is not enough to predict feature mapping, one must also consider how well the specific dependency structures of the modifier and head fit with respect to the mapped property.

Inductive Inference

All else being equal, an immutable feature is more likely than a mutable one to be projected from one object to another. For example, suppose you have just learned that your computer's central processor depends for its operation on built-in memory registers. We suspect that this fact would substantially increase your belief that the central processor of the next computer you come across will also require memory registers. We also suspect that if you have just learned that your computer has a built-in cache, you will be less certain whether to generalize that property to the next computer you see. You already know that a large number of a computer's properties depend on its central processor, so if the central processor depends on memory registers, memory registers must be central and therefore immutable features of your computer. A natural inference is that they are also immutable of other computers. But the news that a computer has a feature, like a cache, without any indication that anything depends on it, seems to provide less sanction for the inductive projection of that feature.

Explanation Generation

Mutability also helps to determine feature relevance when people are constructing explanations. Explanations tend to focus on mutable features (Kahneman & Miller, 1986). For example, appealing to an immutable feature (like *to move*) to explain why some birds have webbed feet is less satisfying than appealing to a mutable property (like *to swim*). Because it is more mutable, *can swim* has a greater chance of not holding, which makes it useful for distinguishing those birds that have webbed feet from those that do not. Of course, a good explanation depends on attributes unrelated to mutability, like the most significant difference between an object and its contextually relevant contrast set. Our claim is only that explainers will be biased to appeal to mutable differences.

Event Understanding

Kahneman and Miller (1986) and Kahneman and Varey (1990) point out that knowledge about mutability is at the heart of our understanding of events. Responses to outcomes are influenced by other outcomes that could have occurred. The reaction to the outcome of a game or contest, for instance, depends on how close the loser came to winning. Whether or not we notice a person's behavior depends directly on our expectations of that person and of people in general. Events are not perceived in isolation but within a background of counterfactuals—events that might have, but did not, occur. This background can determine