Finally, no consistent relation obtained between Equation (1) and the measures of diagnosticity or salience. Some of the individual category correlations for these measures have fairly high magnitude. However, some are positive and others negative, with the result that the average across categories is low. Mutability ratings are the only ones that were always in one direction, a direction that we predicted in each case.

To help interpret the correlations between mutability measures and Equation (1), we compared them to four other models: one representing the centrality of a feature as the sum of dependencies on it; another was based on Equation (1) with an added nonlinearity and a free parameter; another had a term representing the critical feature's dependence on other features; and, finally, we tested a model representing centrality as a feature's total connectivity, without regard for the direction of dependencies. The only model to fit mutability judgments closer than Equation (1) was its nonlinear counterpart requiring a free parameter, and it did only marginally better. The models and their fits are reported in Appendix B.

These analyses provide evidence in support of the hypothesis that features are immutable to the extent they are central in a concept's dependency structure. In particular, the data support the predicted asymmetry: Features are central to the extent that other features depend on them.

Study 3: Varying Dependency Type

The strongest claim embodied by our centrality hypothesis is that mutability is insensitive to the type of relation binding features. Equation (1) presumes only one kind of relation, a generic asymmetric dependency that varies in strength. In other words, we have supposed that relations can be treated as a single type for the sake of combining them to determine mutability. We test this assumption using artificial categories because varying the type of relations embodied by familiar concepts entails dubious claims about how features depend on others.

Study 3a

Using fictitious disease categories, Ahn and Lassaline (in preparation) have varied the dependency structure of symptoms. Figure 2 displays 3 symptoms associated with the disease Yorva. The symptoms are not independent; rather, symptom S causes symptom V which, in turn, causes symptom L. Ahn and Lassaline taught people about such diseases and then asked them to judge the likelihood that someone had disease Yorva if either they displayed symptoms S and V, but not L, or symptoms V and L, but not S. (No figure was presented to the participants in this study.) The results showed that people assigned higher likelihood judgments in the case that L (the "resultant" symptom) was missing than in the case that S (the "causal" symptom) was missing. They judged causes to be less mutable than effects when evaluating disease likelihood.

Our hypothesis is that this result does not occur because of the causal relations per se, but rather because of the asymmetric dependency structure of the diseases. We predict that similar results can be achieved using other kinds of non-causal dependency relations. To