Supplementary Materials: Assumed Correlation Structure under Assumption of Independent Competing Risks.

We use a competing risks duration model in our analysis. Given the various independence assumptions we make (e.g., IIA), there are no restrictions on the range of the correlations we can estimate (either in theory or practice) in a single-risk model. In practice, the same holds true for the correlations in the separately estimated competing risks model. However, a theoretical problem does arise in the competing risks model that, while not actually constraining the range of possible estimates, could lead (for some ranges of these estimates) to a situation in which they cannot be properly interpreted as correlations (implying a theoretical covariance matrix that is not semi-definite). The intuition behind the problem is straightforward: When variables A and B are highly correlated and variables A and C are highly correlated, it must be the case that B and C are also correlated. Correspondingly, if one assumes B and C are independent, then that assumption necessarily limits the correlations between A and B and/or A and C that are consistent with the independence assumption. Applying this to our case, since in the independent competing risk framework we assume that the latent duration until replacement for a potential government j is uncorrelated with its latent duration until dissolution, this implies certain theoretical restrictions on the correlations between these two latent durations and the underlying latent utility for choosing potential government j (or, equivalently, if we estimate both these correlations to be sufficiently large in our separate estimations—which nothing in the model actually prevents—then we will be faced with an interpretation problem). As a helpful reviewer pointed out, this theoretical limit (given the other assumptions in our model) is that the sum of the squares of our two correlations must be less than 1. If this is not the case, then the model estimates would be inconsistent with our assumption of independence between the failure risks, suggesting that a model that relaxes this assumption would be necessary. Gordon (2002) has provided such a model, and in future work it would be useful to combine that model with a government selection component as we do here. As a first step, however, we proceed by assuming that the two competing risks are independent.

As it turns out (see Table 3), our estimate of the correlation coefficient for the replacement risk is 0.310; for the dissolution risk, it is 0.112. Since $0.310^2 + 0.112^2 < 1$, these estimates are consistent with our assumption of independent risks (though again, there is nothing in our model constraining these estimates to meet this restriction).