C The Diebold-Gunther-Tay method

Diebold, Gunther, and Tay (1998) propose a way to evaluating density forecasts. The basic idea is that since under the null hypothesis the forecasts are equal to the true densities (conditioned on past information), applying the cumulative distribution function (the probability integral transform or PIT) to the series of observations should yield a series of iid uniform-$[0, 1]$ variables. Whether the transformed variables are iid uniform can be checked in various ways. Diebold, Gunther, and Tay (1998) suggest plotting histograms and autocorrelation functions to visualize the quality of the density forecasts.

In order to apply the PIT to our predicted recovery rate densities, we create a vector $y^\dagger$ in which we stack all recovery rate observations. For each element in the vector, we can now create a conditional density forecast from our estimated model.

Applying the cumulative distribution function associated with these density forecasts to the vector $y^\dagger$ yields a vector of transformed variables. Under the null hypothesis that the density forecasts are correct, the elements of the vector of transformed variables should be an iid uniform series. Serial correlation of the series would indicate that we have not correctly conditioned on the relevant information. A departure from uniformity would indicate that the marginal distributions are inappropriate.

D Supplementary tables

Table 5
Recovery Rate Statistics by Year

This table reports some annual statistics for the data used in the paper. First column figures are issuer-weighted default rates of US bond issuers provided by Moody's. The other three columns are the number of default events, and the mean and standard deviation for recovery rates in the Altman data.

Year	Default rate	Number of observations	Mean Recovery	Standard Deviation
1981	0.17%	1	12.00	-
1982	1.08%	12	39.51	14.90
1983	1.02%	5	48.93	23.53
1984	0.98%	11	48.81	17.38
1985	1.01%	16	45.41	21.87
1986	2.07%	24	36.09	18.82
1987	1.65%	20	53.36	26.94
1988	1.52%	30	36.57	17.97
1989	2.43%	41	43.46	28.78
1990	4.14%	76	25.24	22.28
1991	3.55%	95	40.05	26.09
1992	1.85%	35	54.45	23.38
1993	1.13%	21	37.54	20.11
1994	0.80%	14	45.54	20.46
1995	1.25%	25	42.90	25.25
1996	0.77%	19	41.90	24.68
1997	0.89%	25	53.46	25.53
1998	1.60%	34	41.10	24.56
1999	2.61%	102	28.99	20.40
2000	3.43%	120	27.51	23.36
2001	4.98%	157	23.34	17.87
2002	3.33%	112	30.03	17.18
2003	2.36%	57	37.33	23.98
2004	1.28%	39	47.81	24.10
2005	1.12%	33	58.63	23.46

Table 6
Recovery Rates by Seniority and Industry

Panel A: Number of observations and the mean and standard deviation of recovery rates in our sample classified by seniority, for the whole sample (all default events), for default events for which we only observe recovery on a single instrument (with only one seniority), and for default events for which we observe recoveries on at least two different seniorities.
Panel B: Number of observations and the mean and standard deviation of recovery rates in our sample classified by industry.

Panel A: Recovery Rates by Seniority			
Seniority	Number of observations	Mean Recovery	Standard Deviation
All default events			
Senior Secured	203	42.08	25.48
Senior Unsecured	366	36.88	23.29
Senior Subordinated	326	32.90	23.77
Subordinated	154	34.51	23.05
Discount	75	21.29	18.48
Default events with single recovery			
Senior Secured	145	39.29	23.35
Senior Unsecured	239	36.45	22.45
Senior Subordinated	209	34.52	23.30
Subordinated	87	37.86	20.22
Discount	29	21.72	19.67
Default events with multiple recoveries			
Senior Secured	58	49.04	29.25
Senior Unsecured	127	37.68	24.87
Senior Subordinated	117	30.00	24.42
Subordinated	67	30.16	25.79
Discount	46	21.03	17.91
Panel B: Recovery Rates by Industry			
Industry			
Building	10	33.56	36.24
Consumer	149	35.66	22.21
Energy	47	36.47	16.66
Financial	95	35.60	25.54
Leisure	69	41.43	29.40
Manufacturing	395	35.08	23.83
Mining	14	35.52	17.50
Services	65	34.16	28.09
Telecom	169	29.43	20.90
Transportation	66	38.07	23.79
Utility	23	51.34	27.97
Others	22	37.94	19.30

E Supplementary Figures

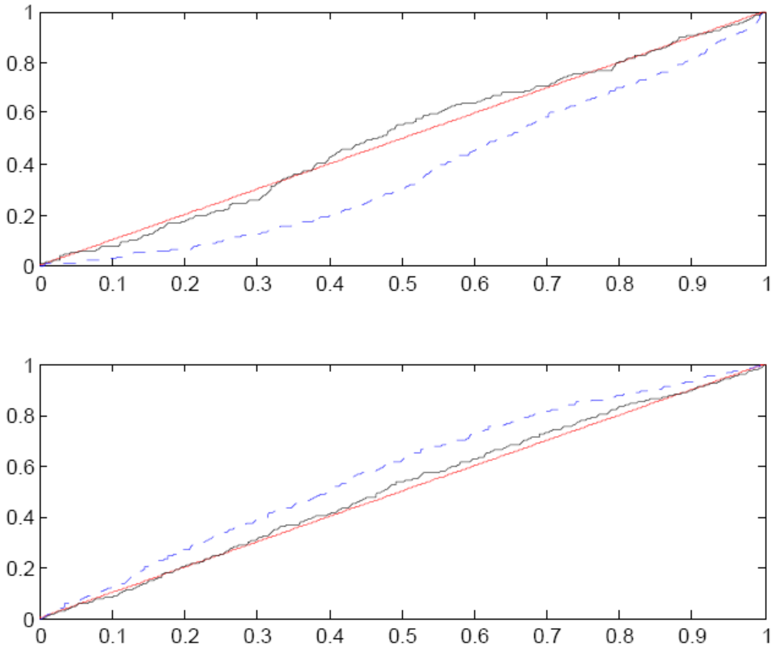

Figure 4. Q-Q plots of the PIT series of the static and cycle models. Dashed line is the static model (M1), solid line is a dynamic model (M2). The upper panel is a Q-Q plot for periods which are identified as upturn by the dynamic model, the lower panel is a Q-Q plot for periods which are identified as downturns by the dynamic model.

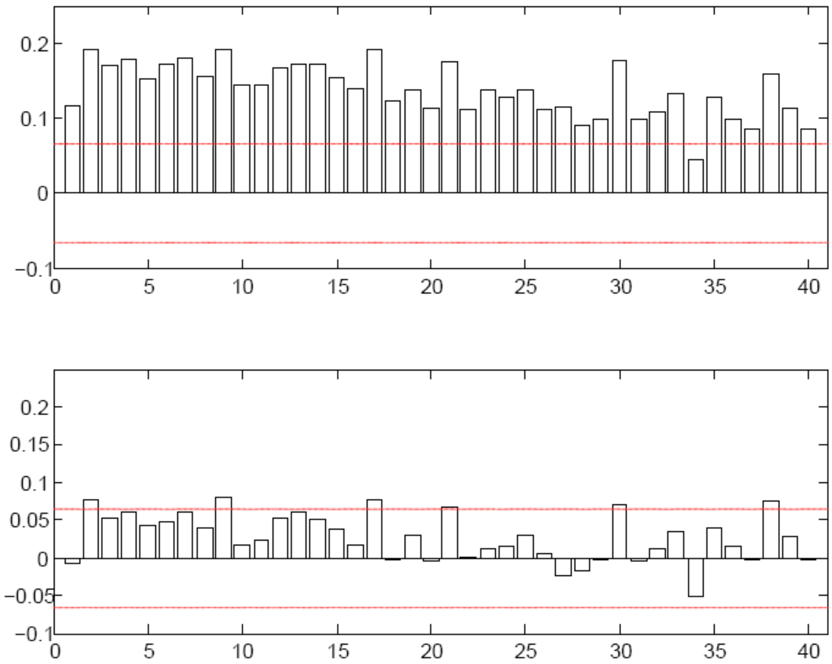

Figure 5. Correlograms of the PIT series of the static and cycle model.

Upper panel is the correlogram of the static model (M1), the lower panel is the correlogram of a dynamic model (M2). Horizontal lines are 5% two-sided confidence intervals for a single autocorrelation coefficient.

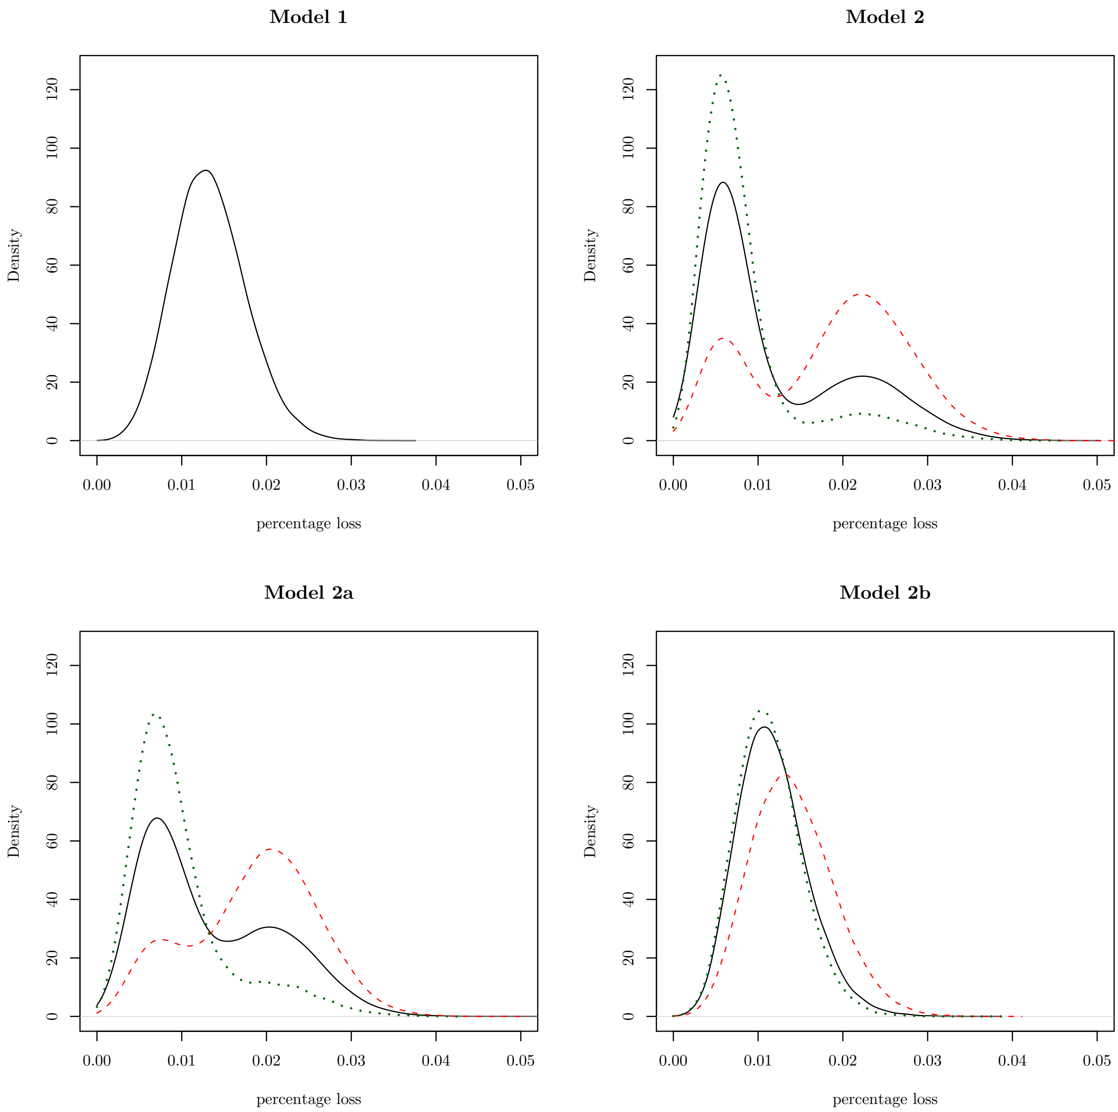

Figure 6. Loss densities for Models 1, 2, 2a, and 2b
Loss densities for Models 1, 2, 2a and 2b. The solid black line is the unconditional loss density (i.e. assuming the probability of being in an upturn is equal to the unconditional probability), dotted green line is the upturn loss density (probability of being in upturn of 1), the red dashed line is the downturn loss density (probability of being in upturn of 0).

References

Diebold, F. X., T. A. Gunther, and A. S. Tay, 1998, "Evaluating Density Forecasts with Applications to Financial Risk Management," *International Economic Review*, 39, 863–883.