

Online Appendix for “Does Unusual News Forecast Market Stress?”

Paul Glasserman and Harry Mamaysky *

September 8, 2018

This Online Appendix covers the following items: a subperiod analysis excluding the financial crisis; details of the impulse responses in the full-sample aggregate VARs from the paper; a robustness check replacing the VIX with a medium-term volatility index; a robustness check using nonparametric impulse response functions; additional details for the panel VAR in the paper; additional details for the VAR with market-specific news; and a VAR analysis incorporating a measure of the volatility risk premium.

I Subperiod analysis

Here we analyze the performance of our single-name panels and the aggregate VAR in the pre-crisis subsample (up to and including December 2006). We first study the single-name panels; and then analyze the aggregate-level VARs.

A Single name panels

We repeated the panel analysis of Section IV from the paper over the pre-crisis time-period from June 1998 to December 2006. We estimated a modified version of equation (8) where we dropped the implied volatility data series (since they only start in January of 2005). We ran

*Glasserman: Columbia Business School and Office of Financial Research (OFR), pg20@columbia.edu. Mamaysky: Columbia Business School, hm2646@columbia.edu.

the panel using 2 and 3 lags, to generate tables comparable to Table 8 in the paper. Tables 1 and 2 show the results.

The results for *SENTPOS* and *ENTPOS* were weaker in the pre-crisis subsample (in fact, the coefficients on *ENTPOS* and *ENSENT_POS* were positive). However the loadings on *ENTNEG*, *SENTNEG* and *ENTSENT_NEG*, while smaller in magnitude, were generally significant and had the appropriate (positive) sign. The stronger results evident in Table 8 suggest that our model was particularly effective in forecasting single name realized volatility during the financial crisis, but negative entropy and sentiment were significant both statistically and economically in forecasting realized volatility even in the pre-crisis period.

B Aggregate impulse responses

We re-estimated the aggregate-level impulse response functions from Section V.B of the paper in the pre-crisis time period from April 1998 to December 2006. Figures 1 (interacted sentiment first) and 2 (sentiment first) show the impulse responses for *ENTSENT_NEG* and *SENTNEG* shocks. Figures 3 (interacted sentiment first) and 4 (sentiment first) show the impulse responses for the positive sentiment shocks, *ENTSENT_POS* and *SENTPOS*.

All aggregate level responses were quantitatively and qualitatively similar to the full-sample results, though the significance levels of the impulse responses were lower (which is at least partially attributable to the shorter data sample) and *SENTPOS* tended to outperform *ENSENT_POS* (*ENTSENT_NEG* still outperformed *SENTNEG*). We conclude that the relationships among S&P realized volatility, the VIX and our aggregate level entropy and sentiment measures were largely the same in the pre- and post-crisis subsamples of the data.

II Full-sample aggregate impulse responses

Here we show details of the impulse response functions (IRFs) in the aggregate VAR which was described in Section V.B of the paper. Figure 5 shows the impulse responses to negative sentiment shocks. The top panel shows the IRFs with *ENTSENT* precedes sentiment in the ordering, and the bottom panel shows the IRFs when *ENTSENT* follows sentiment. Figure 6 shows the analogous IRFs for positive sentiment shocks.

A Including VXMT in the aggregate VAR

It is possible that the hump-shaped responses in our VARs are due to the fact that the news innovations in our sample are systematically about events that will take place four or five months in the future, and the VIX, since it only measures one-month ahead volatility, doesn't react to such news right away. To control for this possibility we rerun our VARs, but include the Mid-Term VIX (ticker VXMT) as the seventh variable (placed between the VIX and *SPX_RVOL*). The VXMT is constructed using S&P 500 options with 6-to-9 months left to expiration, and provides a six-month ahead volatility forecast. Though we have less data for this augmented VAR, the shape of the impulse responses of the VIX, VXMT and *SPX_RVOL* to news innovations are qualitatively similar to our original specification. For example, the maximal response of the VXMT and of the VIX to an *ENTSENT_NEG* innovation occurs in month four (*SPX_RVOL* peaks in month three). This suggests that the market does not fully incorporate all relevant news about future volatility into the VXMT price.

The VXMT is the medium-term VIX index with option maturities between 6 and 9 months. Our data on the VXMT start in January 2008 (the VAR excluding VXMT runs from April 1998 to December 2015). From January 2008 to July 2016 the correlation between the VXMT and the 6 month 90% strike S&P 500 implied volatility (obtained from Bloomberg) is 99.52%. This S&P 500 implied volatility series starts in January 2005, so we use that as our synthetic VMXT series, in order to have more data for our VAR. Figure 7 shows plots of the two volatility series. Figures 8 and 9 show the impulse responses (using a Cholesky decomposition of the covariance matrix) to *ENTSENT_NEG* and *ENTSENT_POS* innovations when the order of the VAR variables is: *VIX*, *VXMT*, *SPX_rvol*, *ENTSENT_NEG*, *SENTNEG*, *ENTSENT_POS*, and *SENTPOS*.

Running the VAR with either the 90% strike implied volatilities or with the VXMT series – but run over the shorter time window from starting in 2008 – yields almost identical results.

III Nonparametric impulse response functions

Our impulse response functions are based on a linear VAR model. As a robustness check on the patterns we observe, we estimated nonparameteric IRFs using an event study approach. We do the following for both volatility measures (implied and realized) and all four news measures:

- Find months in which volatility is in the middle 40% of its empirical distribution over the full sample;
- Within those months, find months in which NEWS_MEASURE is in the top or bottom 20% of its empirical distribution. Interpret these as months of typical volatility that experience a shock to NEWS_MEASURE;
- Trace the path of volatility over the next 10 months;
- Average the paths that start from high levels of NEWS_MEASURE, average the paths that start from low levels of NEWS_MEASURE, and take the difference of the averages. This is an empirical IRF, showing the historical response of volatility to a shock in NEWS_MEASURE. It does not rely on any assumptions of linearity.

Figure 10 shows the response of the VIX (the four left panels) and S&P500 realized volatility (the four right panels) to observed shocks in the four news measures. The title above each chart shows which news measure applies and the number of paths averaged. For example, the chart in the upper left shows the change in the VIX following a large value of ENTSENT_NEG, starting from 14 months in which ENTSENT_NEG was in the top 20% of its distribution and 15 months in which it was in the bottom 20% of its distribution.

The overall pattern in the figures is quite similar to the IRFs we estimate using the VAR model in the paper. All responses go in the expected directions (up for negative news measures, down for positive news measures), and the response to ENTSENT_NEG looks slightly more significant than the response to SENTNEG.¹ The responses look somewhat less significant than under the VAR specification, but this is to be expected because the nonparametric method throws out most of the data. The VAR model makes better use of the full data set by imposing additional structure. The nonparametric results in Figure 10 provide a robustness check to the VAR results.

IV Panel VAR: Additional impulse response functions

To contrast our aggregate results with corresponding results at the company-specific level, we estimate a panel VAR. Our state variables are implied volatility, realized volatility, and the news

¹The dashed lines in the figure show ± 2 standard errors away from the mean path. The standard errors do not account for serial dependence and are only approximate.

measures *ENTSENT_NEG*, *SENTNEG*, *ENTSENT_POS*, and *SENTPOS*, all measured at the company-specific level. To estimate the VAR, for each variable we run a panel regression of the form in equation (8) in the main paper) with the right side consisting of two lags of all six variables. Stacking the coefficient vectors from these regressions yields the two coefficient matrices (corresponding to the two lags) for the VAR model, from which we calculate impulse response functions.

We use bootstrapping to estimate confidence bands around each impulse response function, as follows. Each row in our dataset corresponds to a company-month pair and records values of the six state variables for that company in that month. We have 3513 rows of complete data. To generate a bootstrap sample, we randomly resample 3513 rows (with replacement) from the original data. We estimate the panel VAR for the bootstrap sample and then calculate impulse response functions. We repeat this for 100 bootstrap samples. The confidence bands show the range of the middle 95% of impulse response functions at each time step.

Figure 11 shows IRFs for shocks to *ENTSENT_NEG* (top)/*ENTSENT_POS* (bottom) and *SENTNEG* (top)/*SENTPOS* (bottom) where *ENTSENT_NEG* and *ENTSENT_POS* precede *SENTNEG* and *SENTPOS* in the VAR ordering. Figure 12 shows IRFs for shocks to *ENTSENT_NEG* and *SENTNEG* where the ordering is reversed.

V VAR with market-specific news

Figure 13 shows the distribution of article lengths (by total word count) for the roughly 38,000 market-specific articles chosen based on the filter described in the paper. We then considered different subsets of these articles based on word length (for example, restricting ourselves to articles between 100 and 500 words yields just under 20,000 articles) but found this had no qualitative impact on the resulting sentiment and entropy series, or on the VAR results; therefore we report only results for the entire set of market-specific articles. Figure 14 shows the number of articles per month that our filter generates, as well as the positive and negative sentiment, and aggregate entropy, series. There is a spike in articles during the financial crisis. The data run from April 1998 to December 2014, for a total of 201 months, and 191 articles per month on average. In the post crisis period the average number of articles per month falls to 90. In contrast to our bottoms-up series, the market-specific positive and negative sentiment series are positively correlated.

Figure 15 shows the six series that enter the VAR described in the body of the paper, with the VIX (and their correlations with the VIX) superimposed for reference. These market-specific sentiment series are characterized by: (i) spikes around the financial crisis; (ii) a positive correlation between the positive and negative sentiment series (see Figure 14, and note that our bottom-up series are negatively correlated as shown in Table 10 in the paper); and (iii) a positive (zero) correlation between *ES_NEG_MKT* (*ES_POS_MKT*) and the VIX. Figure 16 shows the impulse responses from the VAR model.

VI Volatility Risk Premium

As a rough measure of a volatility risk premium (VRP) at time t , we use the difference between the VIX in month t and S&P 500 realized volatility in month $t + 1$, interpreting the VIX as the risk-neutral expectation of realized volatility in the next month. (More precisely, VIX^2 is the risk-neutral expectation of realized variance; the results are not qualitatively different if we use variance instead of volatility.)

Figure 17 shows the response of our VRP measure to shocks in *ENTSENT_NEG* and *SENTNEG* for our aggregate VAR. As in the paper, we run two versions of the VAR with two different orderings of the variables. The figure shows (particularly the first and last panels) that an increase in negative news produces a drop in the VRP, which gradually mean-reverts. This is consistent with our observation in the paper that volatility does not immediately fully incorporate the information in the news measures. In addition, Bollerslev et al. (2009) document that higher values of the VRP forecast higher returns, so the drop we see in the VRP is consistent with negative news forecasting lower returns. In Figure 18, we see that a shock to positive news produces an increase in the VRP (consistent with the opposite effect for negative news) but the effect is not significant.

Figure 19 shows corresponding results for our single-name panel VAR. The drop in the VRP following a shock to *ENTSENT_NEG* is now more instantaneous, consistent with the quicker response of realized volatility in the single-name IRFs in the paper. The response to positive news in Figure 20 is positive but not significant. Except for the speed of the response to negative news, the results are qualitatively similar in the single-name and aggregate settings.

VII Incremental R^2 of text measures for future volatility

The incremental explanatory power of our news measures is in the same range as that of variables used to forecast the equity risk premium. In one-month ahead forecasting regressions for aggregate market returns, in-sample R^2 's range from 0% to 2% across most forecasting variables (such as divided-price ratio, earnings-price ratio, T-bill rate, term spread, default spread, etc.) with the median R^2 below 1% (see Fama and French (1988), Campbell and Thompson (2008), and Welch and Goyal (2008)).² Campbell and Thompson (2008) argue that variables with such low in-sample R^2 's will have even lower out-of-sample R^2 's but importantly remain economically significant for investors despite this (see the discussion on page 1525 of Campbell and Thompson 2008).

To evaluate the incremental explanatory power of our text measures, we examine the drop in R^2 that occurs in our month-ahead forecasting regression, equation (10) in the paper, when removing all our text measures from the right-hand side of the regression, in both the aggregate and panel versions of equation (10).

Removing the news measures from the set of explanatory variables drops the R^2 of the single-name regression forecasting implied volatility by 0.005 (row labeled “ivol s-n” in Table 3), and of the realized volatility regression by 0.007 (row labeled “rvol s-n”). For comparison, the incremental R^2 's from lagged implied volatility range from 0.025 to 0.027 and the incremental R^2 from lagged realized volatility range from 0.041 to 0.045. In the aggregate version of the paper's equation (10), the incremental R^2 of the text measures is 0.030 in the implied S&P 500 volatility forecasting regression (row labeled “ivol macro”) and 0.036 in the S&P 500 realized volatility regressions (row labeled “rvol macro”). The incremental R^2 for lagged implied volatility ranges from 0.071 to 0.094, and for lagged realized volatility the incremental R^2 ranges from 0.004 to 0.012. In short, the incremental R^2 's from our news measures are meaningful when compared with the incremental R^2 from the best predictors of volatility, and their magnitude is in line with the explanatory power of variables that have been shown to forecast the aggregate equity risk premium. The latter are considered to be economically useful predictors of returns despite their low in-sample R^2 's (as argued in Campbell and Thompson 2008).

²The consumption-to-wealth ratio is the outlier with an R^2 of 4.57% according to Campbell and Thompson (2008) though only 1.88% according to Welch and Goyal (2008) in monthly forecasting regressions.

References

- Campbell, J. and S. Thompson, “Predicting excess returns out of sample: Can anything beat the historical average?” *Review of Financial Studies*, 21 (2008), 1509–1531.
- Fama, E. and K. French, “Dividend yields and expected stock returns,” *Journal of Financial Economics*, 22 (1988), 3–25.
- Welch, I. and A. Goyal, “A comprehensive look at the empirical performance of equity premium prediction,” *Review of Financial Studies*, 21 (2008), 1455–1508.

VIII Tables and figures

Single name panel with 2 lags in pre-crisis period

	rvol-only (1)	rvol-only (2)	rvol-only (3)	rvol-only (4)	rvol-only (5)	rvol-only (6)	rvol-only (7)
rvol_l1	0.372***	0.376***	0.398***	0.397***	0.375***	0.370***	0.369***
rvol_l2	0.209***	0.211***	0.215***	0.219***	0.207***	0.218***	0.215***
ret_mi_l1	0.802***	0.867***	0.812***	0.811***	0.865***	0.862***	0.856***
ret_mi_l2	0.134***	0.109**	0.058	0.062	0.101**	0.127**	0.123**
ARTPERC_l1	-0.538**	-0.833***	-0.684**	-0.655**	-0.868***	-0.757***	-0.747***
ARTPERC_l2	-0.003	0.055	0.192	0.132	0.011	0.046	0.025
SENTPOS_l1		-0.396**					
SENTPOS_l2		0.240					
ENTPOS_l1			0.474**				
ENTPOS_l2			0.097				
ENTSENT_POS_l1				-0.083			
ENTSENT_POS_l2				0.525***			
SENTNEG_l1					0.131		
SENTNEG_l2					0.362*		
ENTNEG_l1						0.139	
ENTNEG_l2						0.049	
ENTSENT_NEG_l1							0.189
ENTSENT_NEG_l2							0.253
Sum Last 2	-0.54**	-0.156	0.571**	0.442*	0.493*	0.188	0.442
Sum Last 2 p-val	[0.035]	[0.557]	[0.049]	[0.082]	[0.088]	[0.491]	[0.137]
R2 adj	0.447	0.470	0.490	0.490	0.470	0.482	0.483
Start date	Apr 1996	Jun 1998	Jun 1998	Jun 1998	Jun 1998	Jun 1998	Jun 1998
End date	Dec 2006	Dec 2006	Dec 2006	Dec 2006	Dec 2006	Dec 2006	Dec 2006

Table 1: This table reports the results of the single name panel model. The dependent variable is shown in the column heading, with the regressors in the rows. The row labeled “Sum Last Two” shows the sum of the two bottom-most coefficients in each column. The regression is run with individual fixed effects. Residuals are clustered by time for computing standard errors. ‘*’, ‘**’, and ‘***’ indicate significance at the 0.10, 0.05, and 0.01 levels respectively. Dates over which each panel is run are shown in the “Start” and “End” date rows.

Single name panel with 3 lags in pre-crisis period

	rvol-only (1)	rvol-only (2)	rvol-only (3)	rvol-only (4)	rvol-only (5)	rvol-only (6)	rvol-only (7)
rvol_l1	0.333***	0.340***	0.374***	0.376***	0.339***	0.339***	0.339***
rvol_l2	0.125***	0.133***	0.142***	0.148***	0.131***	0.140***	0.137***
rvol_l3	0.165***	0.155***	0.129***	0.130***	0.153***	0.137***	0.138***
ret_mi_l1	0.808***	0.855***	0.780***	0.773***	0.855***	0.857***	0.850***
ret_mi_l2	0.221***	0.193***	0.158***	0.161***	0.185***	0.236***	0.230***
ret_mi_l3	-0.033	-0.038	-0.036	-0.039	-0.044	-0.018	-0.030
ARTPERC_l1	-0.511**	-0.717***	-0.562*	-0.554*	-0.749***	-0.672**	-0.701**
ARTPERC_l2	-0.060	0.024	0.110	0.034	0.007	-0.010	0.000
ARTPERC_l3	0.101	0.068	0.045	0.073	0.044	0.037	0.046
SENTPOS_l1		-0.415**					
SENTPOS_l2		0.280					
SENTPOS_l3		0.008					
ENTPOS_l1			0.409*				
ENTPOS_l2			-0.024				
ENTPOS_l3			0.491**				
ENTSENT_POS_l1				-0.115			
ENTSENT_POS_l2				0.485**			
ENTSENT_POS_l3				-0.077			
SENTNEG_l1					0.078		
SENTNEG_l2					0.269		
SENTNEG_l3					0.066		
ENTNEG_l1						-0.001	
ENTNEG_l2						-0.093	
ENTNEG_l3						0.777***	
ENTSENT_NEG_l1							0.159
ENTSENT_NEG_l2							0.069
ENTSENT_NEG_l3							0.443**
Sum Last 3	-0.47*	-0.127	0.876**	0.293	0.413	0.683**	0.671*
Sum Last 3 p-val	[0.066]	[0.686]	[0.030]	[0.360]	[0.201]	[0.038]	[0.054]
R2 adj	0.463	0.484	0.494	0.493	0.484	0.493	0.492
Start date	May 1996	Jul 1998	Jul 1998	Jul 1998	Jul 1998	Jul 1998	Jul 1998
End date	Dec 2006	Dec 2006	Dec 2006	Dec 2006	Dec 2006	Dec 2006	Dec 2006

Table 2: This table reports the results of the single name panel model. The dependent variable is shown in the column heading, with the regressors in the rows. The row labeled “Sum Last Two” shows the sum of the two bottom-most coefficients in each column. The regression is run with individual fixed effects. Residuals are clustered by time for computing standard errors. ‘*’, ‘**’, and ‘***’ indicate significance at the 0.10, 0.05, and 0.01 levels respectively. Dates over which each panel is run are shown in the “Start” and “End” date rows.

Text based indicator contribution to volatility forecasting

	LHS	full_R2	start	end	ivol	rvol	news
ivol s-n	ivol	0.582	Mar 2005	Dec 2014	0.027	0.041	0.005
ivol s-n F-test					0.007	0.000	0.008
rvol s-n	rvol	0.576	Mar 2005	Dec 2014	0.025	0.045	0.007
rvol s-n F-test					0.000	0.000	0.006
ivol macro	ivol	0.757	Jun 1998	Dec 2014	0.094	0.004	0.030
ivol macro F-test					0.000	0.235	0.000
rvol macro	rvol	0.615	Jun 1998	Dec 2014	0.071	0.012	0.036
rvol macro F-test					0.000	0.077	0.048

Table 3: The *Full model* column gives unadjusted R^2 's from the regression in (10). The columns labeled *ivol*, *rvol*, and *news* show the drop in the full model unadjusted R^2 if *ivol*, *rvol* or the 4 news measures are removed from the right side of equation (10). The rows labeled “F-test” show the p-value of an F-test comparing the restricted to the unrestricted model (the standard errors are clustered by time for the panel regressions, and Newey-West with auto-lag selection is used for the macro tests). The top half of the table refers to the single name panel regressions for implied and realized volatilities with the full model coefficient estimates shown in Tables 7 and 8. The lower half of the table shows the results of the VIX (here called *ivol macro*) and the *SPX_RVOL* (*rvol macro*) regressions from the VARs discussed in Section V.

Aggregate VARs in the pre-crisis time period

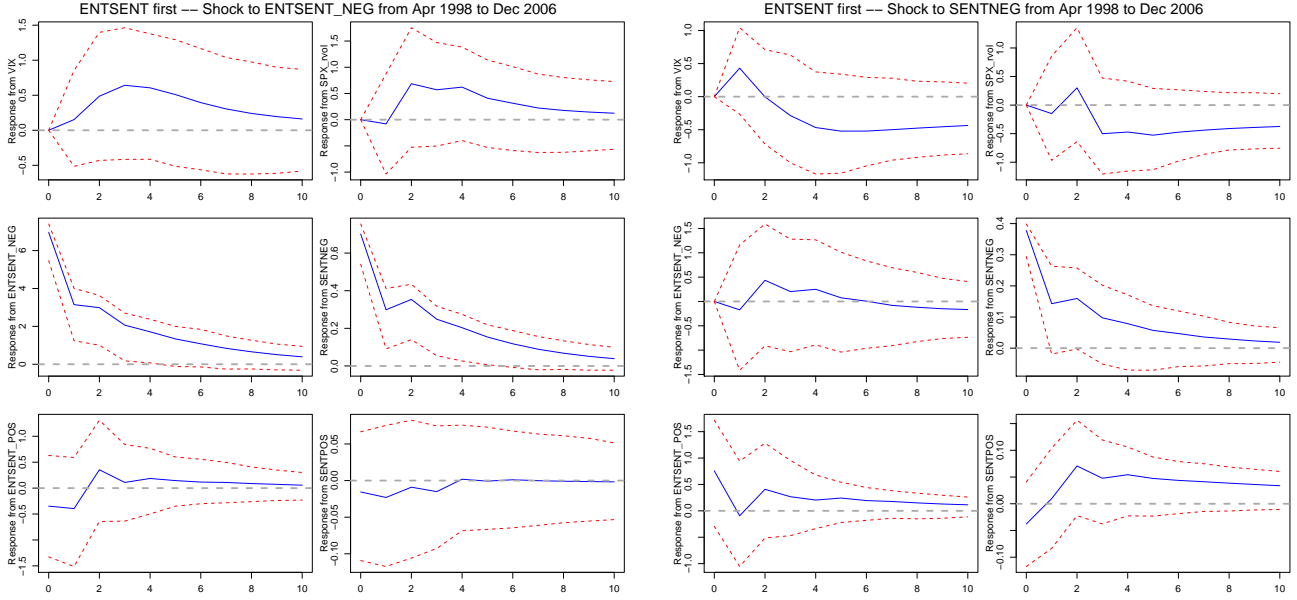


Figure 1: Impulse response functions for a shock to *ENTSENT_NEG* (left) and *SENTNEG* (right). The order of the variables in the VAR model matches the order of the figures in each block of six, reading left to right, then top to bottom. Dashed lines show 95 percent bootstrap confidence intervals. The horizontal time axis is in months.

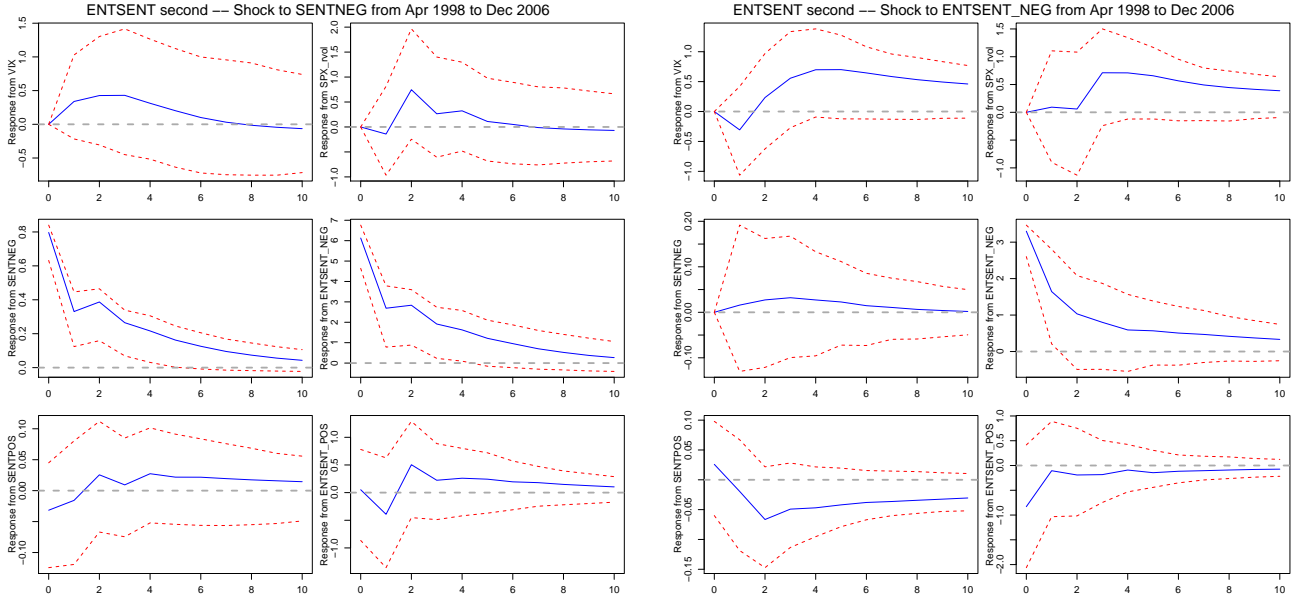


Figure 2: Impulse response functions for a shock to *SENTNEG* (left) and *ENTSENT_NEG* (right). The order of the variables in the VAR model matches the order of the figures in each block of six, reading left to right, then top to bottom. Dashed lines show 95 percent bootstrap confidence intervals. The horizontal time axis is in months.

Aggregate VARs in the pre-crisis time period

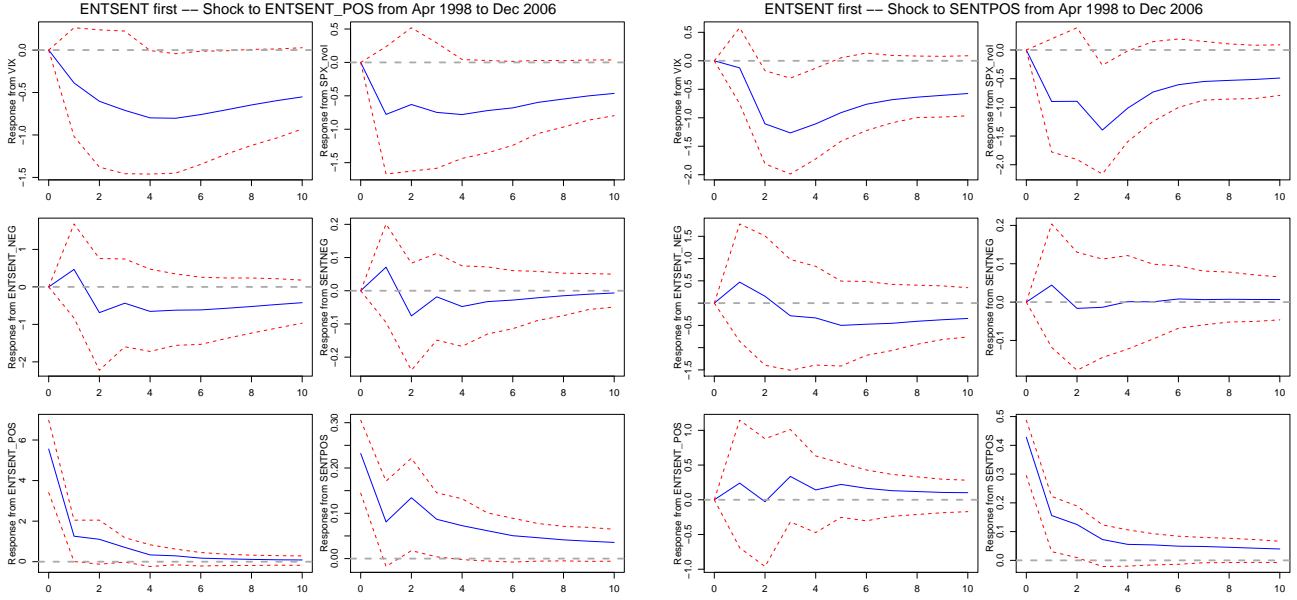


Figure 3: Impulse response functions for a shock to *ENTSENT_POS* (left) and *SENTPOS* (right). The order of the variables in the VAR model matches the order of the figures in each block of six, reading left to right, then top to bottom. Dashed lines show 95 percent bootstrap confidence intervals. The horizontal time axis is in months.

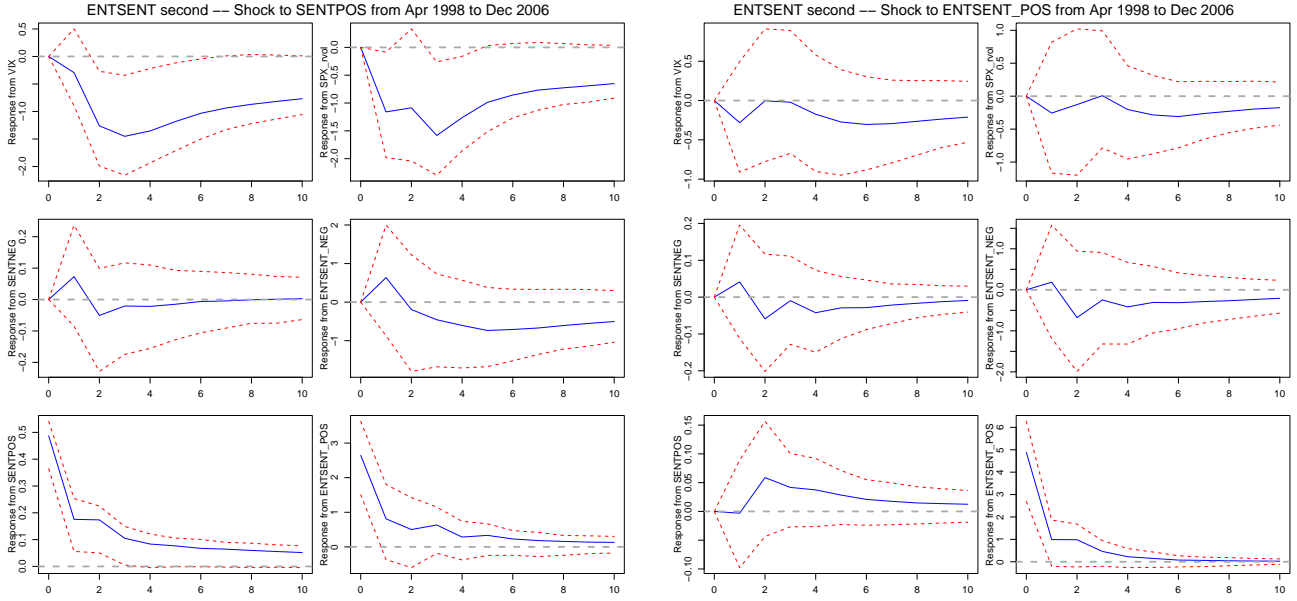
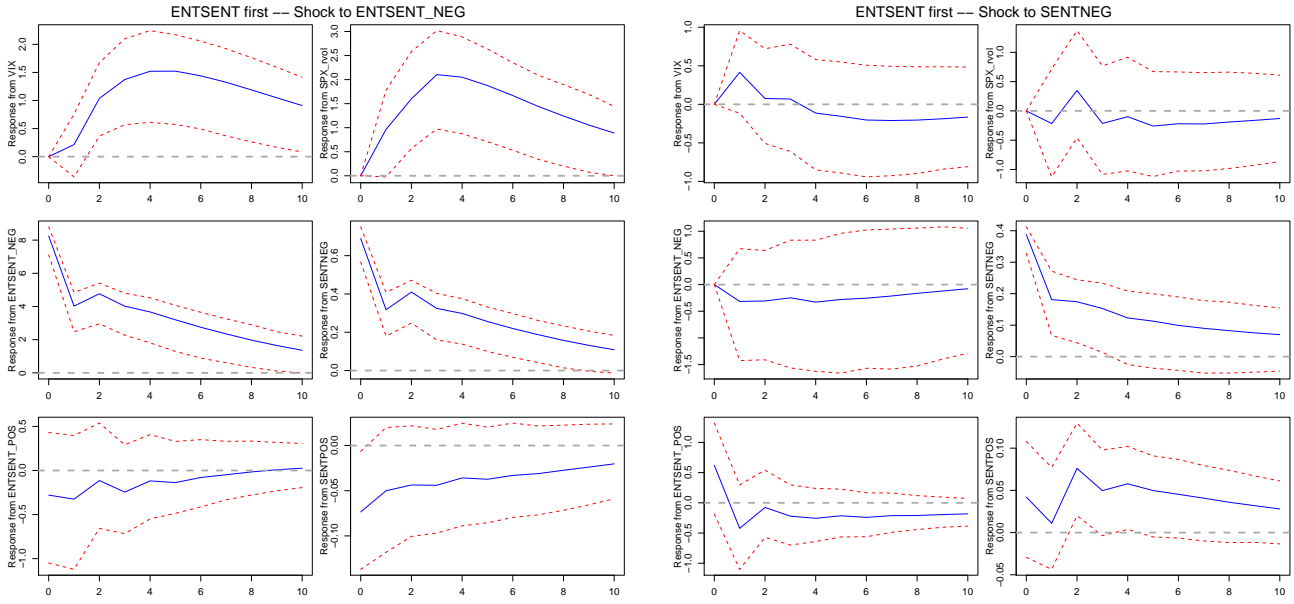


Figure 4: Impulse response functions for a shock to *SENTPOS* (left) and *ENTSENT_POS* (right). The order of the variables in the VAR model matches the order of the figures in each block of six, reading left to right, then top to bottom. Dashed lines show 95 percent bootstrap confidence intervals. The horizontal time axis is in months.

Full-sample aggregate VARs – Impulse responses to negative news

ENTSENT before sentiment in ordering



Sentiment before *ENTSENT* in ordering

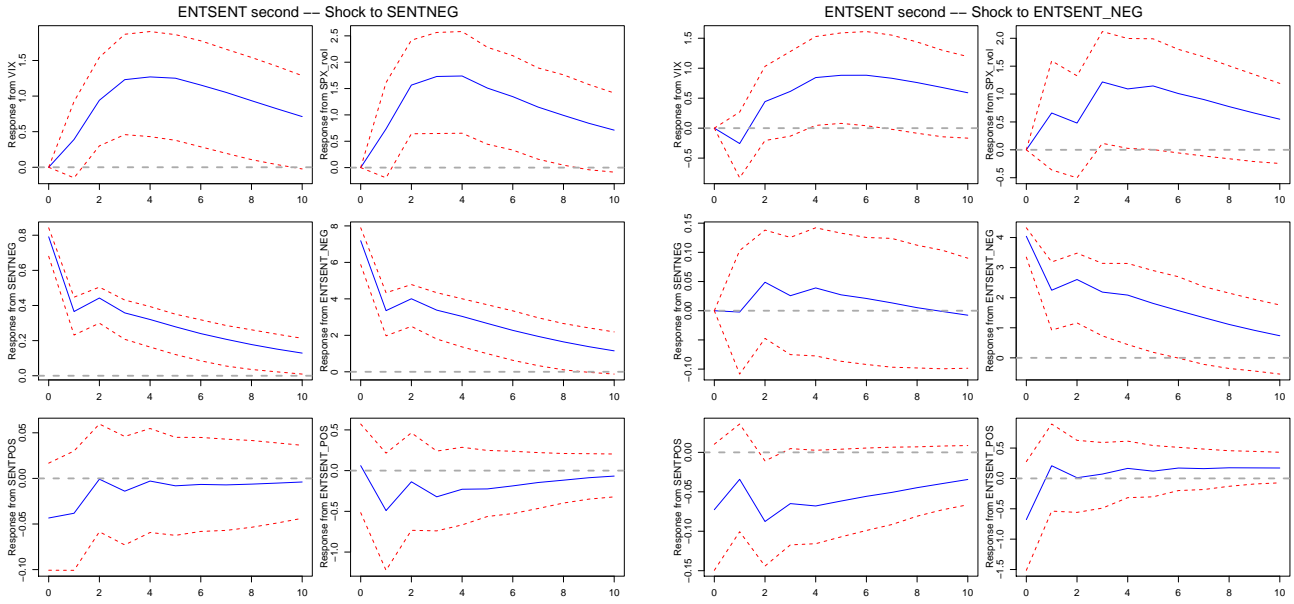
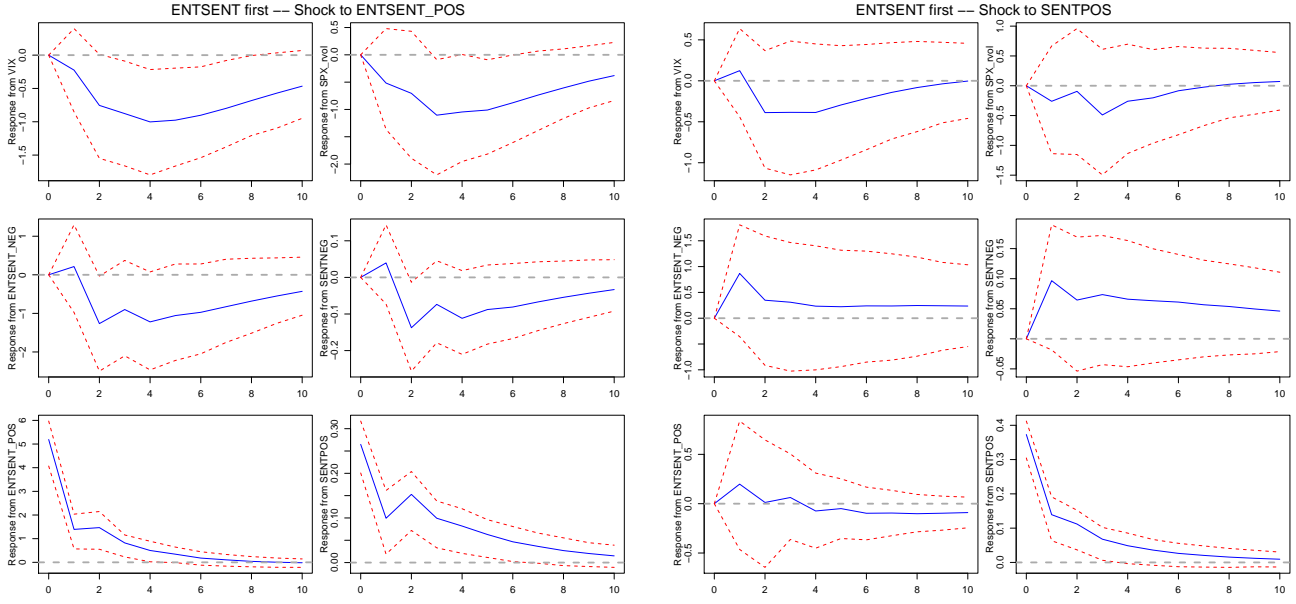


Figure 5: Impulse response functions for a shock to *ENTSENT_NEG* and *SENTNEG*. The order of the variables in the VAR model matches the order of the figures in each block of six, reading left to right, then top to bottom. The top panel ordering has *ENTSENT* before sentiment; the bottom panel has sentiment before *ENTSENT*. Dashed lines show 95 percent bootstrap confidence intervals. The horizontal time axis is in months.

Full-sample aggregate VAR – Impulse responses to positive news

ENTSENT before sentiment in ordering



Sentiment before *ENTSENT* in ordering

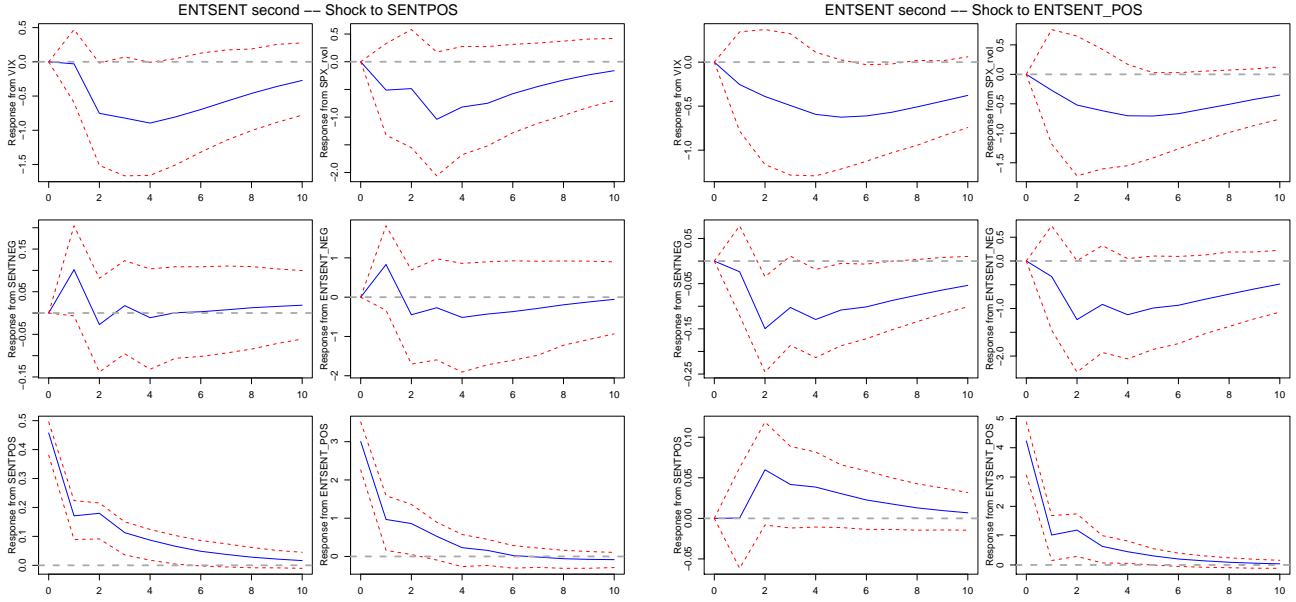


Figure 6: Impulse response functions for a shock to *ENTSENT_POS* and *SENTPOS*. The order of the variables in the VAR model matches the order of the figures in each block of six, reading left to right, then top to bottom. The top panel ordering has *ENTSENT* before sentiment; the bottom panel has sentiment before *ENTSENT*. Dashed lines show 95 percent bootstrap confidence intervals. The horizontal time axis is in months.




Figure 7: VXMT and 90% strike 6 month implied volatility for the S&P 500 (using Bloomberg's volatility surfaces).

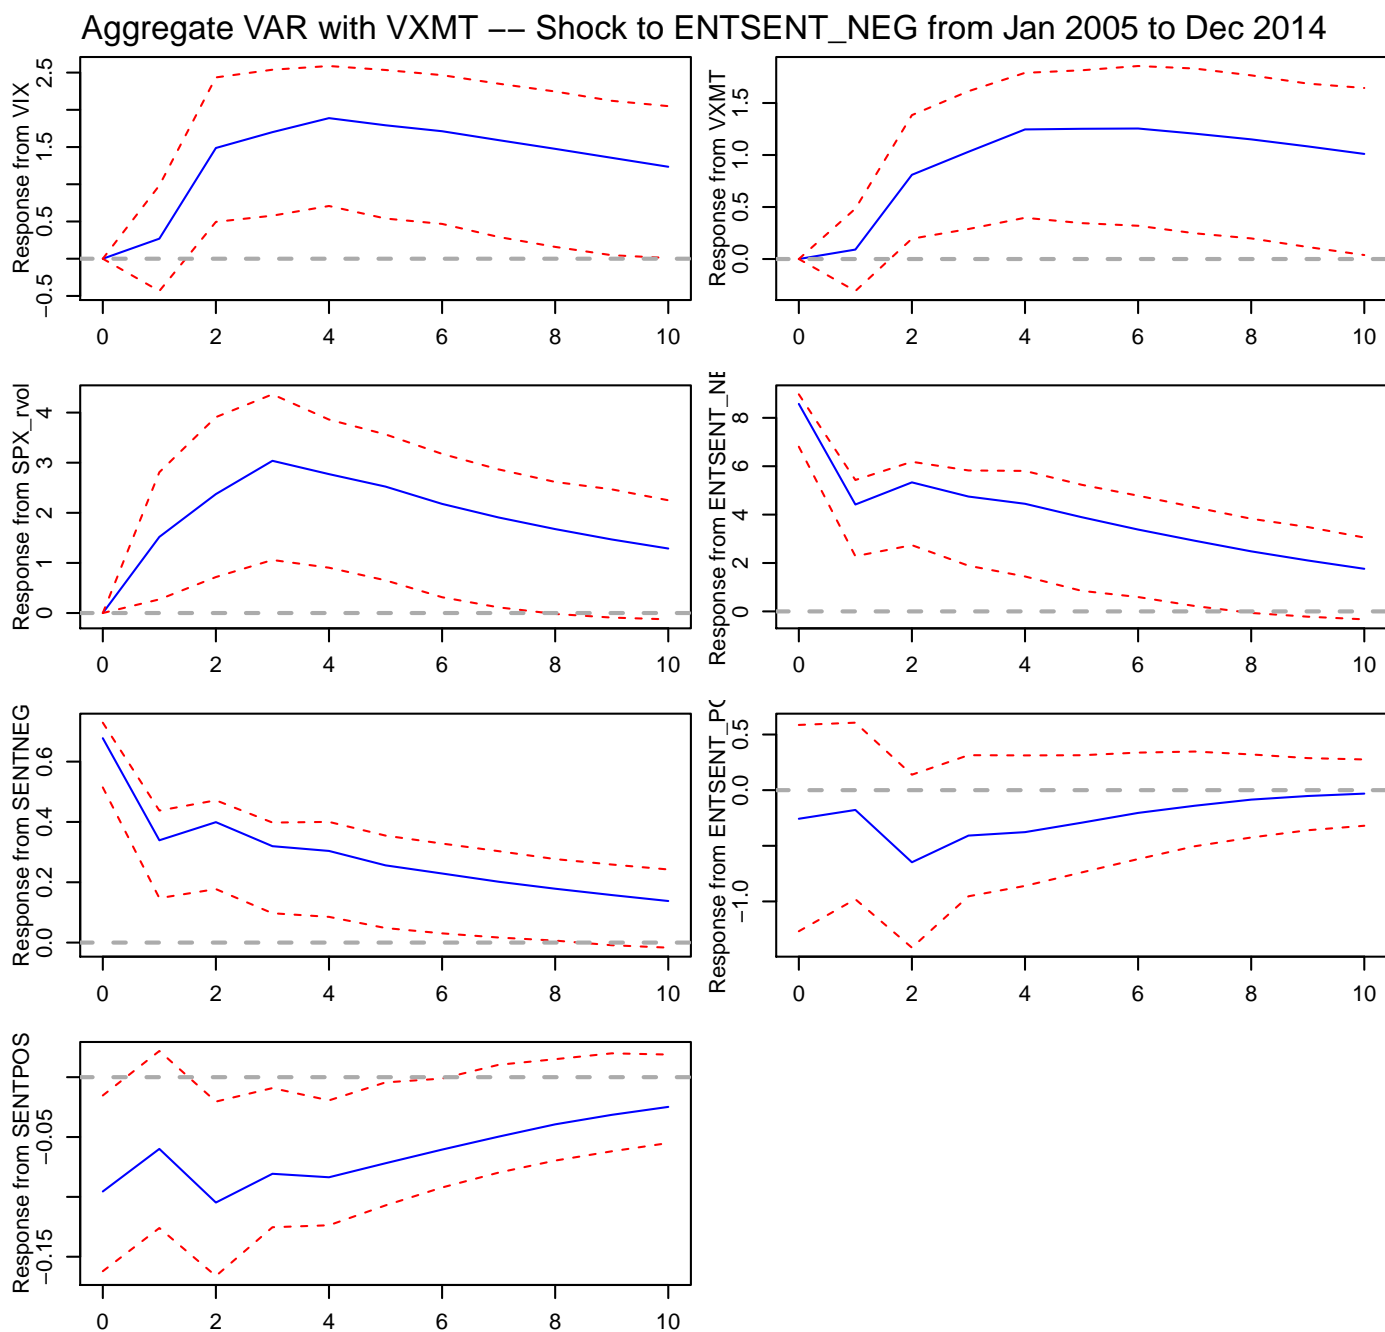


Figure 8: VARs with VXMT.

Aggregate VAR with VXMT --- Shock to ENTSENT_POS from Jan 2005 to Dec 2014

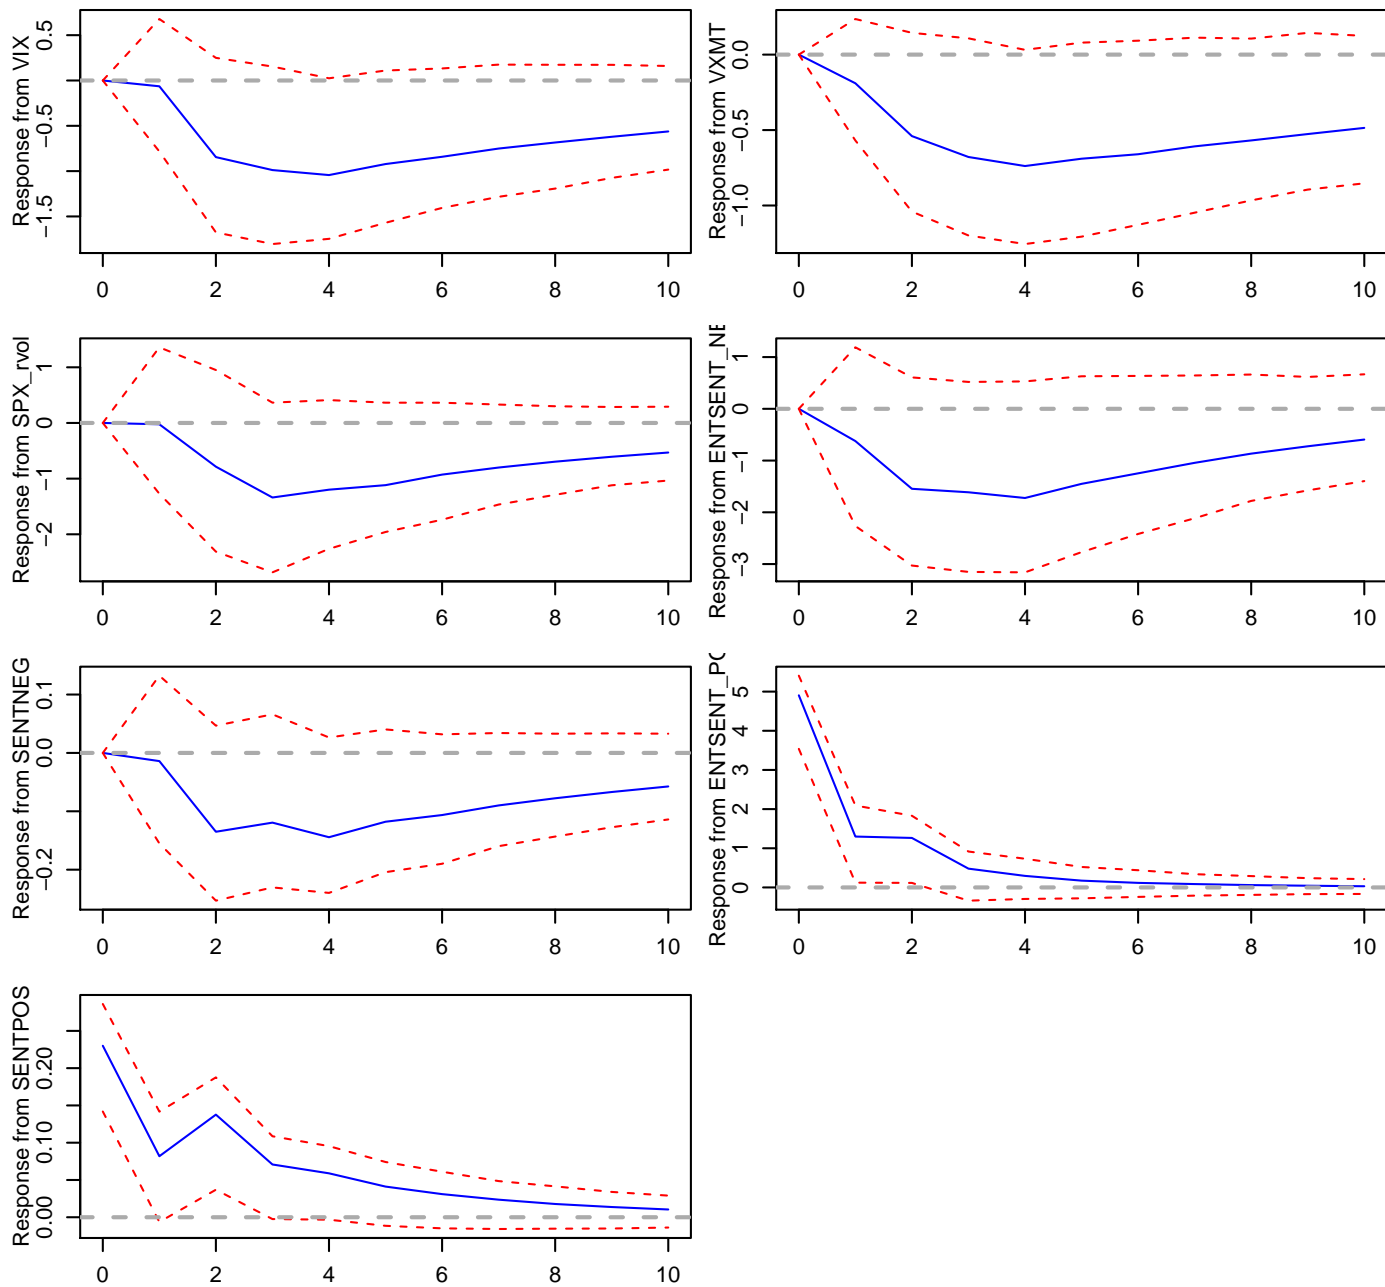


Figure 9: VARs with VXMT.

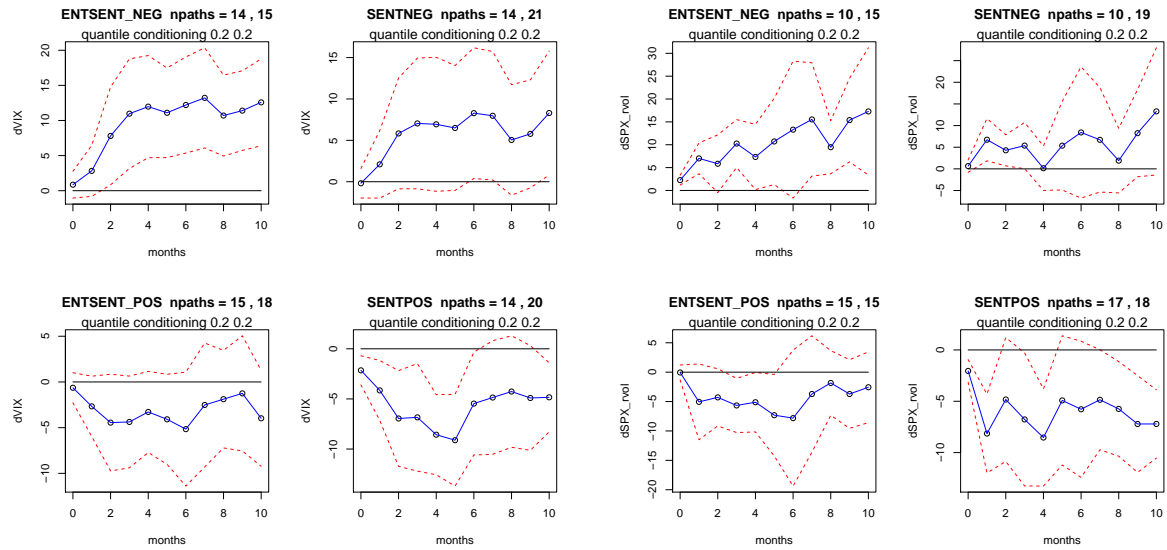


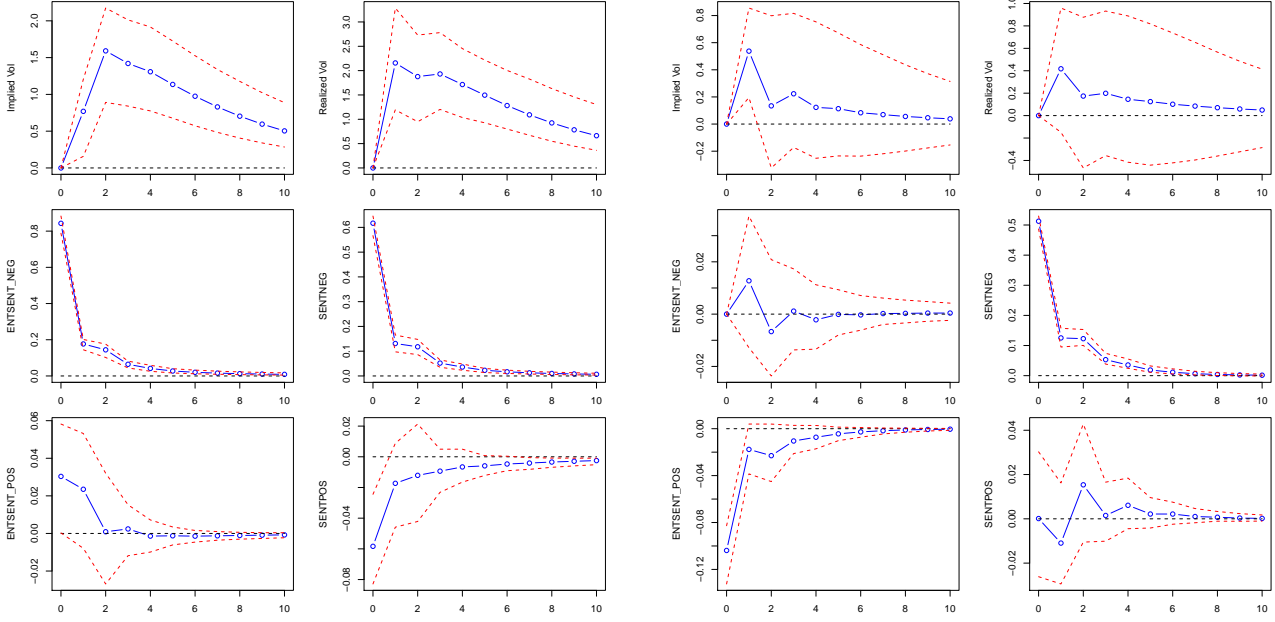
Figure 10: Nonparametric impulse response functions estimated from months in which volatility is in the middle 40% of its distribution and a news measure is in the top or bottom 20% of its distribution. Results shown are for VIX on the left and realized S&P500 volatility on the right. The news measure is shown above each chart.

Company-level panel VAR

Impulse responses to negative news

ENTSENT_NEG shock – *ENTSENT* first

SENTNEG shock – *ENTSENT* first



Impulse responses to positive news

ENTSENT_POS shock – *ENTSENT* first

SENTPOS shock – *ENTSENT* first

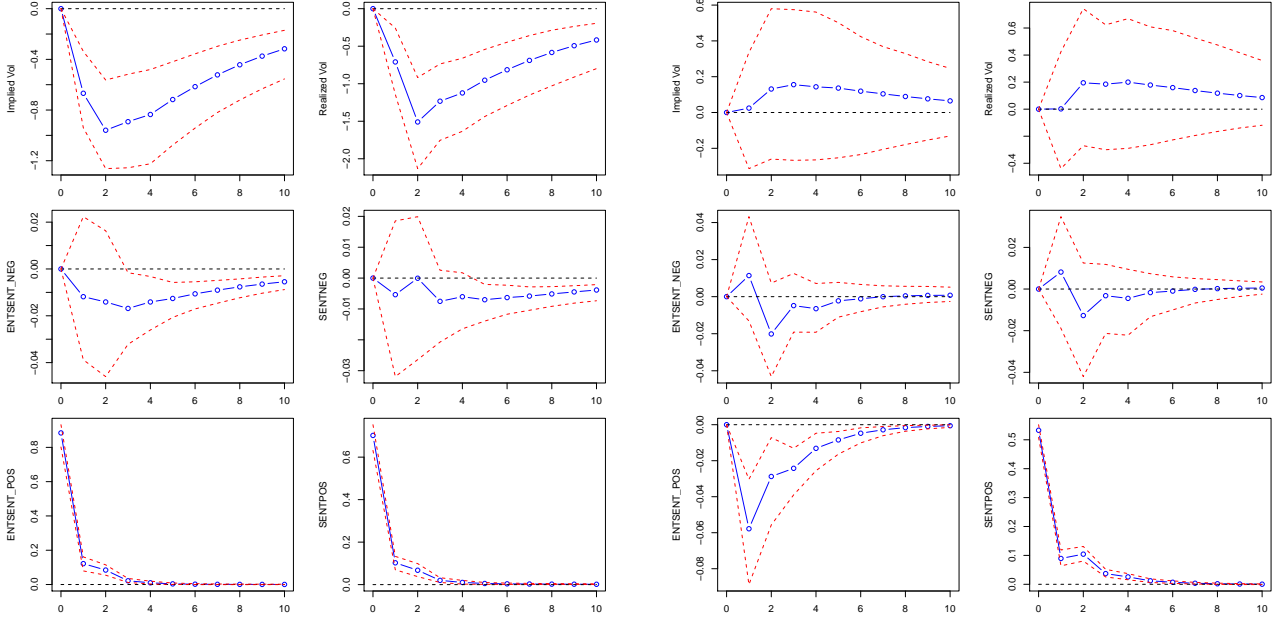


Figure 11: Impulse response functions for a shock to *ENTSENT* (left) and sentiment (right) in the company-level panel VAR. The order of the variables in the VAR model matches the order of the figures in each block of six, reading left to right, then top to bottom. In the figures *ENTSENT* comes before sentiment. Dashed lines show 95 percent bootstrap confidence intervals. The horizontal time axis is in months.

Company-level panel VAR

Impulse responses to negative news

ENTSENT_NEG shock – *SENTNEG* first

SENTNEG shock – *SENTNEG* first

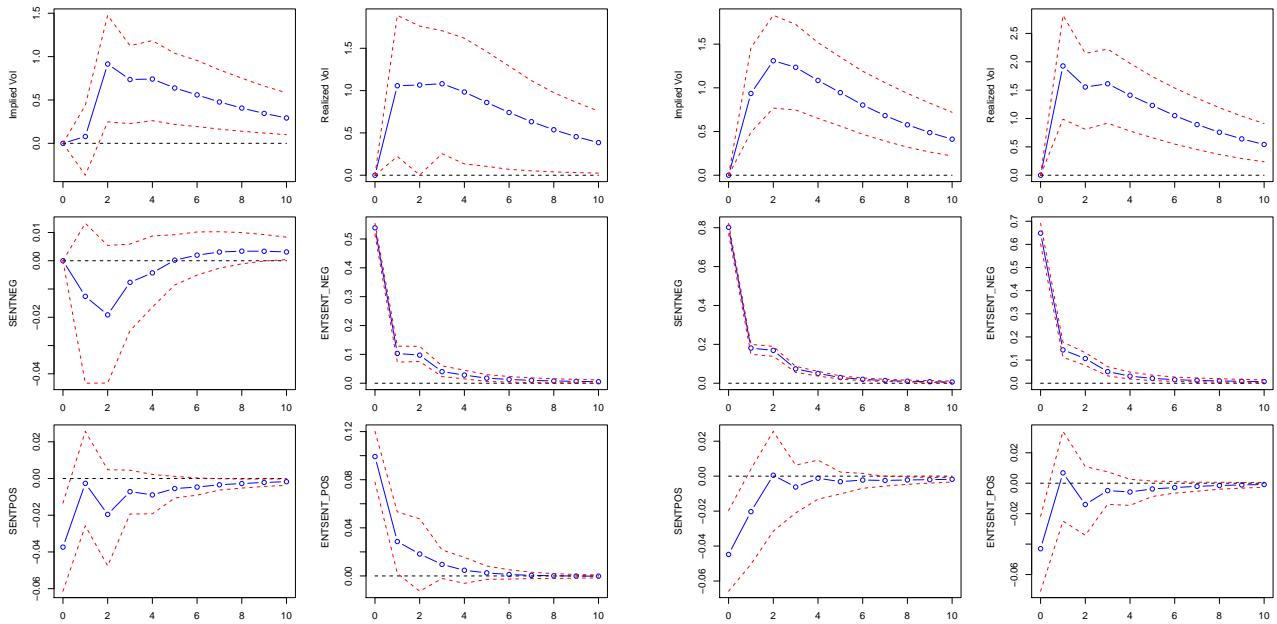


Figure 12: Impulse response functions for a shock to *ENTSENT_NEG* (left) and *SENTNEG* (right) in the company-level panel VAR. The order of the variables in the VAR model matches the order of the figures in each block of six, reading left to right, then top to bottom. Dashed lines show 95 percent bootstrap confidence intervals. The horizontal time axis is in months.




Figure 13: Histogram of words per article for articles satisfying the market-specific news filter from the paper.




Figure 14: Articles counts and market-specific sentiment and entropy over time. The top left chart shows the number of articles per month, on average, in the post-crisis part of the sample (2010 and after).

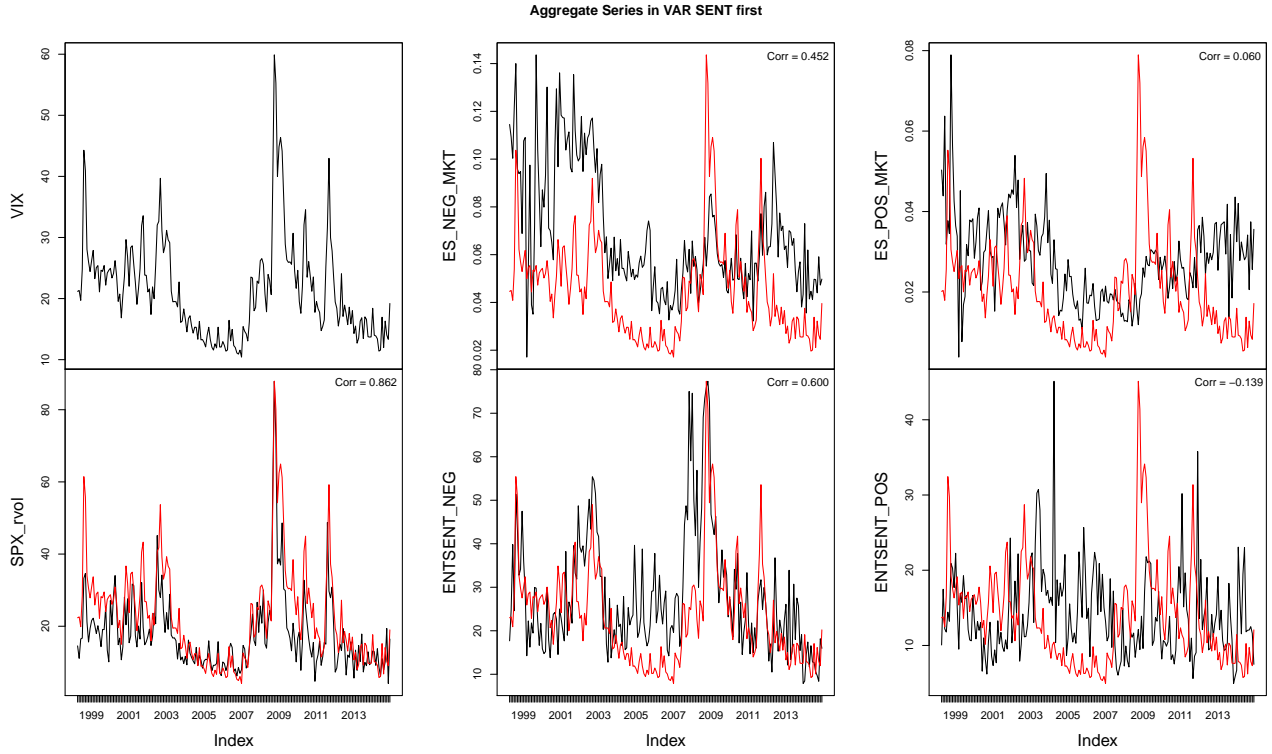
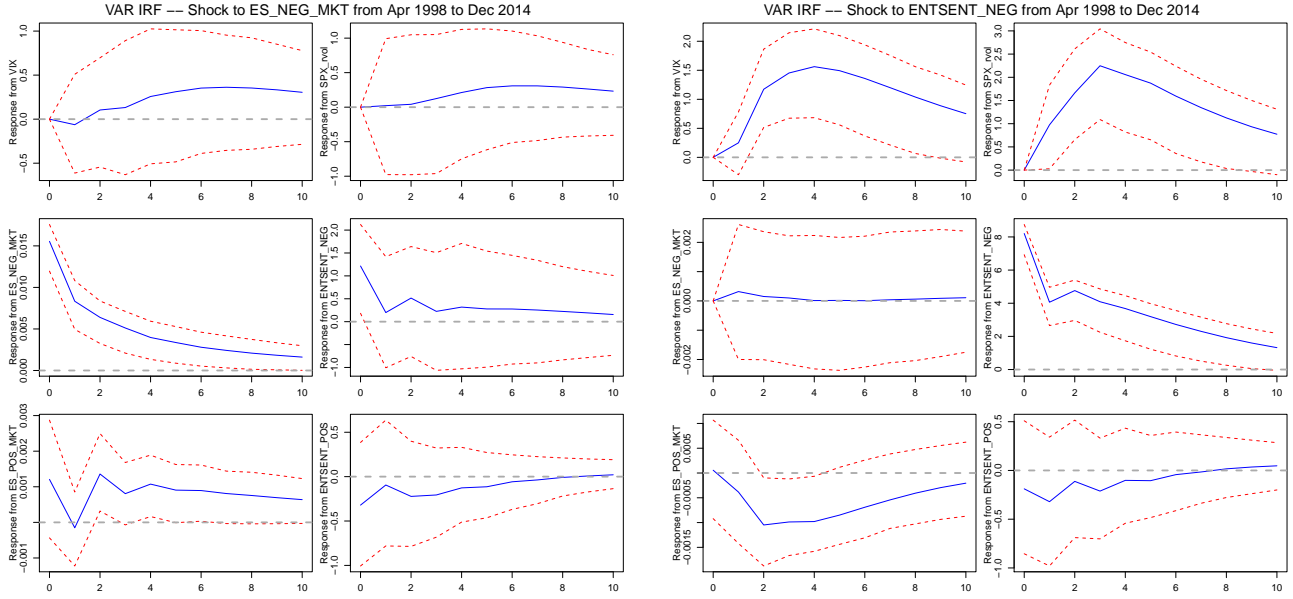


Figure 15: Time series of the components of the aggregate-level VAR employing market-specific sentiment measures in addition to *ENTSENT_NEG* and *ENTSENT_POS*. The VIX index is superimposed (in red) on all the series. The correlation with the VIX is shown in the upper hand corner.

Impulse response from VAR including market-specific news

Negative sentiment shocks



Positive sentiment shocks

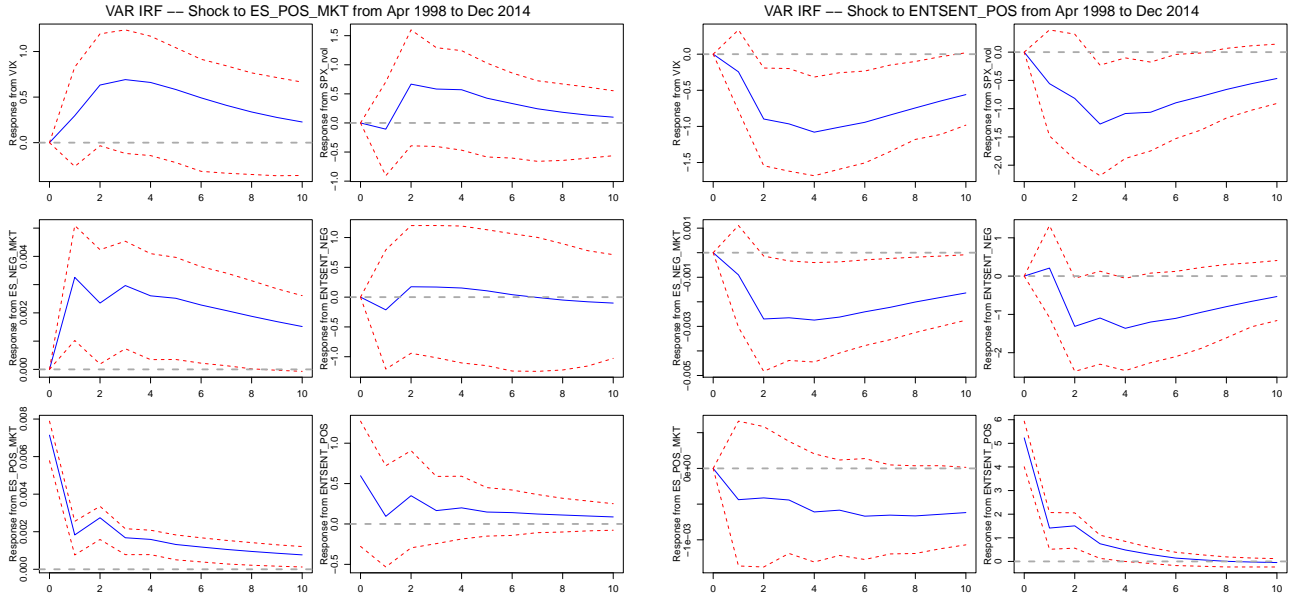


Figure 16: Impulse response functions for a shock to ES_NEG_MKT/ES_POS_MKT (left) and $ENTSENT_NEG/ENTSENT_POS$ (right) in the aggregate-level VAR with 2 lags. The order of the variables in the VAR model matches the order of the figures in each block of six, reading left to right, then top to bottom. Dashed lines show 95 percent bootstrap confidence intervals. The horizontal time axis is in months. Standard errors are bootstrapped using 500 simulation runs.

Aggregate VAR: Impulse response of VRP to negative sentiment shocks

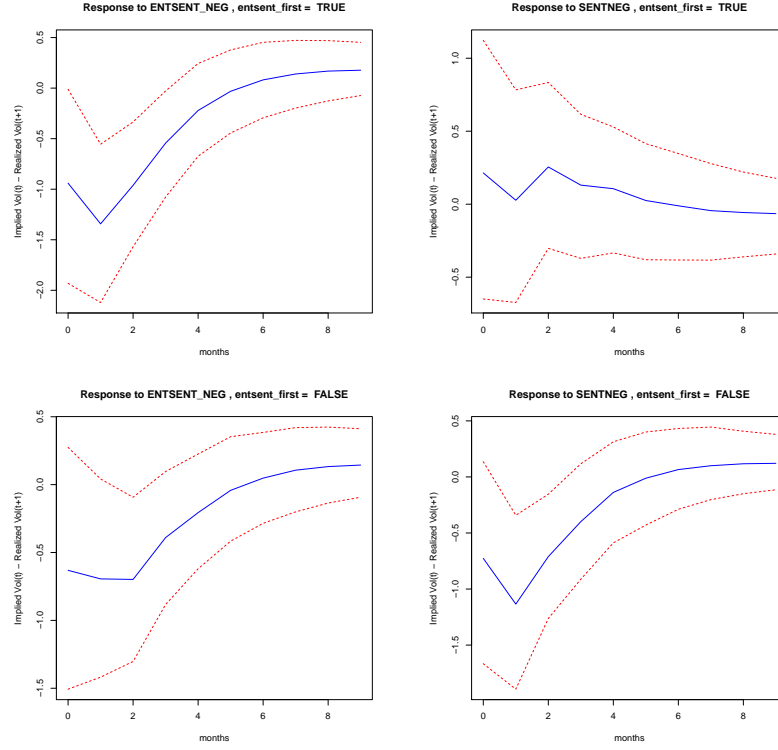


Figure 17: Impulse response functions showing the response of the VRP measure to a shock in *ENTSENT_NEG* (left) or *SENTNEG* (right). The top row corresponds to a VAR specification in which *ENTSENT_NEG* is listed before *SENTNEG*; the bottom row is based on the opposite ordering.

Aggregate VAR: Impulse response of VRP to positive sentiment shocks

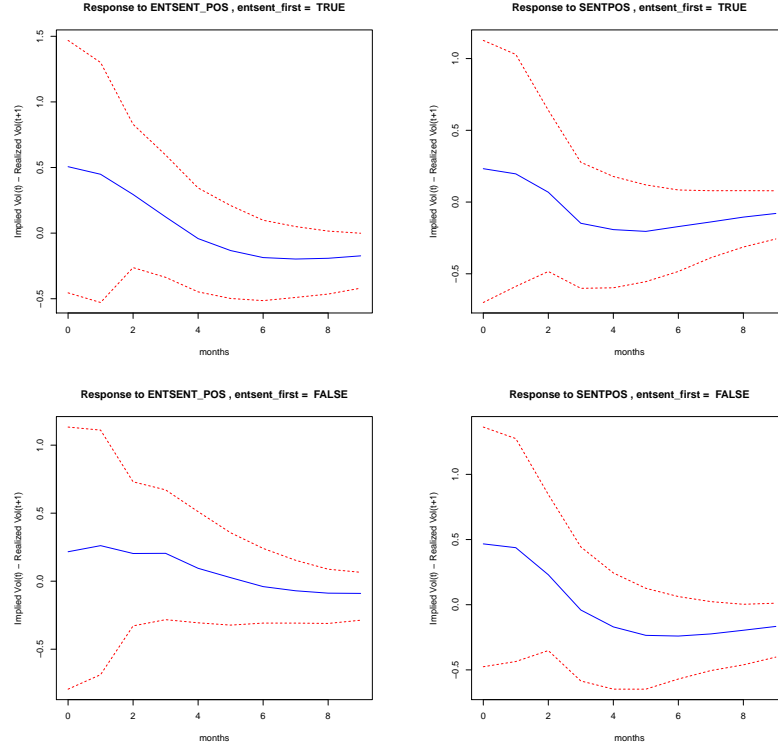


Figure 18: Impulse response functions showing the response of the VRP measure to a shock in *ENTSENT_POS* (left) or *SENTPOS* (right). The top row corresponds to a VAR specification in which *ENTSENT_POS* is listed before *SENTPOS*; the bottom row is based on the opposite ordering.

Panel VAR: Impulse response of VRP to negative sentiment shocks

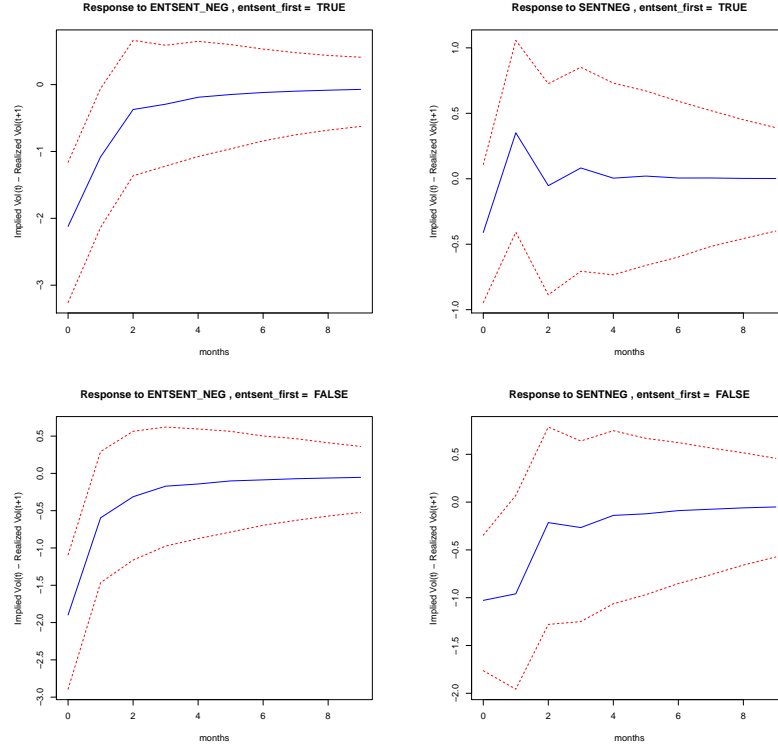


Figure 19: Impulse response functions for panel VAR of individual stocks showing the response of the VRP measure to a shock in *ENTSENT_NEG* (left) or *SENTNEG* (right). The top row corresponds to a VAR specification in which *ENTSENT_NEG* is listed before *SENTNEG*; the bottom row is based on the opposite ordering.

Panel VAR: Impulse response of VRP to positive sentiment shocks

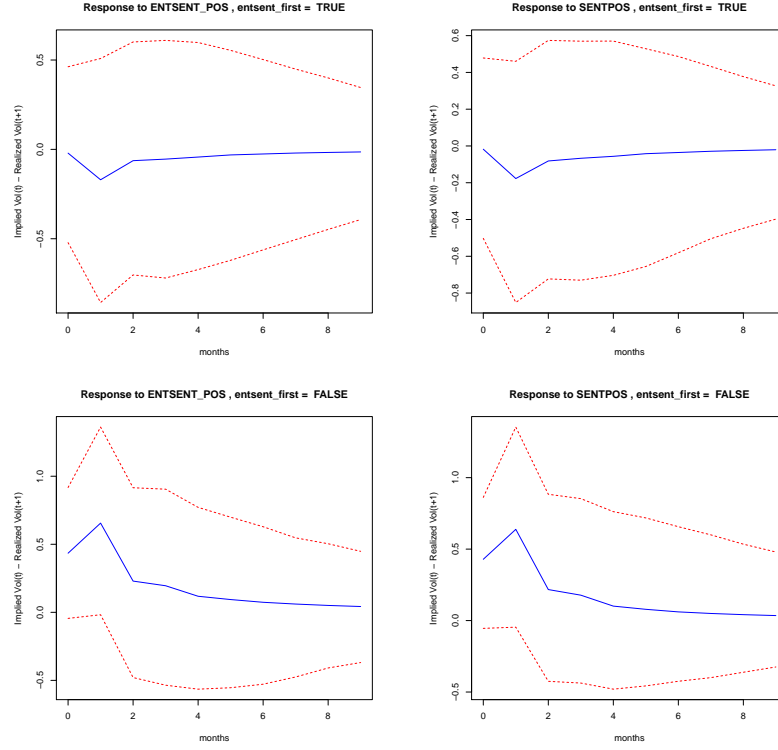


Figure 20: Impulse response functions for panel VAR of individual stocks showing the response of the VRP measure to a shock in *ENTSENT_POS* (left) or *SENTPOS* (right). The top row corresponds to a VAR specification in which *ENTSENT_POS* is listed before *SENTPOS*; the bottom row is based on the opposite ordering.