



**Center for Analytical Finance
University of California, Santa Cruz**

Working Paper No. 3

**Assessing systematic risk in the S&P500 index between
2000 and 2011: A Bayesian nonparametric approach**

Abel Rodríguez, Ziwei Wang and Athanasios Kottas

June 4, 2014

Abstract

We develop a Bayesian nonparametric model to assess the effect of systematic risks on multiple financial markets, and apply it to understand the behavior of the S&P500 sector indexes between January 1, 2000 and December 31, 2011. More than prediction, our main goal is to understand the evolution of systematic and idiosyncratic risks in the U.S. economy over this particular time period, leading to novel sector-specific risk indexes. To accomplish this goal, we model the appearance of extreme losses in each market using a superposition of two Poisson processes, one that corresponds to systematic risks that are shared by all sectors, and one that correspond to the idiosyncratic risk associated with a specific sector. In order to capture changes in the risk structure over time, the intensity functions associated with each of the underlying components are modeled using a Dirichlet process mixture model. Among other interesting results, our analysis of the S&P500 index suggests that there are few idiosyncratic risks associated with the consumer staples sector, whose extreme negative log returns appear to be driven mostly by systematic risks.

About CAFIN

The Center for Analytical Finance (CAFIN) includes a global network of researchers whose aim is to produce cutting edge research with practical applications in the area of finance and financial markets. CAFIN focuses primarily on three critical areas:

- Market Design
- Systemic Risk
- Financial Access

Seed funding for CAFIN has been provided by Dean Sheldon Kamieniecki of the Division of Social Sciences at the University of California, Santa Cruz.

Assessing systematic risk in the S&P500 index between 2000 and 2011: A Bayesian nonparametric approach

Abel Rodríguez, Ziwei Wang and Athanasios Kottas *

Abstract: We develop a Bayesian nonparametric model to assess the effect of systematic risks on multiple financial markets, and apply it to understand the behavior of the S&P500 sector indexes between January 1, 2000 and December 31, 2011. More than prediction, our main goal is to understand the evolution of systematic and idiosyncratic risks in the U.S. economy over this particular time period, leading to novel sector-specific risk indexes. To accomplish this goal, we model the appearance of extreme losses in each market using a superposition of two Poisson processes, one that corresponds to systematic risks that are shared by all sectors, and one that correspond to the idiosyncratic risk associated with a specific sector. In order to capture changes in the risk structure over time, the intensity functions associated with each of the underlying components are modeled using a Dirichlet process mixture model. Among other interesting results, our analysis of the S&P500 index suggests that there are few idiosyncratic risks associated with the consumer staples sector, whose extreme negative log returns appear to be driven mostly by systematic risks.

KEY WORDS: Dirichlet process mixture modeling; Non-homogeneous Poisson process; Non-parametric Bayes; Systematic risk.

*A. Rodríguez is Associate Professor, Department of Applied Mathematics and Statistics, University of California, Santa Cruz, CA; Z. Wang is Statistical Analyst, ask.com, Oakland, CA; and A. Kottas is Professor, Department of Applied Mathematics and Statistics, University of California, Santa Cruz, CA. This research was supported in part by the National Science Foundation under awards SES 1024484 and DMS 0915272, and by funds granted by SIGFIRM at the University of California, Santa Cruz.

1 Introduction

Systematic risk can be loosely defined as the vulnerability of a financial market to events that affect all (or, at least, most) of the agents and products in the market. This is in contrast to *idiosyncratic risks*, which are risks to which only specific agents or products are vulnerable.

The notions of systematic and idiosyncratic risk play a key role in motivating investment diversification. In this paper, we study the evolution of systematic and idiosyncratic risks in the U.S. economy by focusing on the behavior of the S&P500 index and its sector components. The Standard & Poor's 500, or S&P500 index, is a commonly watched stock market index in the U.S., constructed as a market-value weighted average of the prices of the common stock of 500 publicly traded companies. Standard & Poor's, which publishes the index, selects the companies included in the S&P500 index to be representative of the industries in the U.S. economy. These companies are commonly grouped into ten economic sectors: consumer discretionary, consumer staples, energy, financials, health care, industrials, materials, information technology, telecommunication services, and utilities. The largest sector (consumer discretionary) includes 81 companies, whereas the smallest (telecommunication services) includes only 8. In addition to the overall S&P500 index, Standard & Poor's publishes separate indexes for each of these sectors. The behavior of these sector-specific indexes is of independent interest; for example, the performance of the industrial and consumer discretionary components of the S&P500 is used by some analysts as a leading indicator of future economic growth.

The most widely used tool to characterize idiosyncratic risks in financial markets is the Capital Asset Pricing Model (CAPM) (Treynor, 1961, 1962; French, 2003), which uses simple linear regression to relate the expected returns of an individual security or sector to those of the market. The estimated regression coefficient (the so-called "beta") measures the sensitivity of the expected excess asset returns to the expected excess market returns, with larger values of "beta" indicating investments with substantial idiosyncratic risks.

The original CAPM has been repeatedly criticized as being too simplistic (e.g., Fama and French, 2004), and extensions have been driven by considerations both empirical (e.g., the three factor model discussed in Fama and French, 1992) and theoretical (e.g., the behavioral model discussed in Daniel et al., 2001). Two obvious concerns with estimates of idiosyncratic risk based

on the CAPM are the assumption that deviations from the model follow Gaussian distributions, and their reliance on expected returns. Indeed, from a risk management perspective, it is more appealing to define these concepts on the basis of the behavior of extreme returns. As the financial crises of 2007 and 2008 demonstrated, the behavior of markets (e.g., the level of correlation among asset prices) during periods of distress can dramatically deviate from their behavior during periods of calm. Variations of CAPM that focus on extreme returns include Barnes and Hughes (2002), Allen et al. (2009), and Chang et al. (2011). In these papers, quantile regression instead of ordinary linear regression is used to relate the returns of individual securities to those of the market.

This paper develops a novel approach to estimate systematic and idiosyncratic risks across sectors of the U.S. economy that does not involve regression analysis and instead focuses on the relative frequency of extreme losses. More specifically, we model the time of appearance of extreme losses in each market using a superposition of two Poisson processes, one that corresponds to systematic risks that are shared by all sectors, and one that corresponds to the idiosyncratic risk associated with a specific sector. In order to capture changes in the risk structure over time, the intensity functions associated with each of the underlying components are modeled using a Dirichlet process mixture model (Ferguson, 1973; Escobar and West, 1995). Hence, our model can be conceptualized as an example of a Cox process (Cox, 1955). Unlike the CAPM, the proposed methodology does not rely (implicitly or explicitly) on the assumption that returns arise from a Gaussian distribution. Furthermore, our model is dynamic in nature, allowing for the structure of the different risks to evolve over time. As with the CAPM, the main goal of our analysis is a general description of the structure of the different risks and explanation rather than prediction. For example, our model generates idiosyncratic, sector-specific risk indexes, as well as metrics that quantify the overall relative importance of idiosyncratic and systematic risks on each sector.

To motivate the class of models we develop, consider the daily returns associated with the ten sectors making up the S&P500 index (see also Section 4). Figure 1 presents the most extreme *negative* log returns on four of those sectors between January 1, 2000 and December 31, 2011. It is clear from the figure that all sectors present an increased frequency of extreme

losses around periods of distress, such as the so-called “dot com” bubble burst in March of 2000 and the climax of the financial crises in September 2008. However, it is also clear that certain features are particular to specific sectors, such as the increased number of extreme returns in the energy sector in 2004 and 2005. Furthermore, even when the frequency of losses tends to increase significantly for all markets, it is not the case that extreme losses occur in all markets on exactly the same dates.

[Figure 1 about here.]

The rest of the paper is organized as follows. Section 2 develops our modeling approach and discusses its main properties. Section 3 presents a description of the computational approach to inference (with technical details included in the Appendix) along with a careful discussion on subjective prior elicitation. Section 4 considers the analysis of the U.S. market using data from the S&P500 index. Finally, Section 5 concludes with discussion.

2 Modeling approach

We focus on the negative log returns for the ten S&P500 sector indexes

$$x_{i,j} = -100 \times \log \left(\frac{S_{i,j}}{S_{i-1,j}} \right),$$

where $S_{i,j}$ is the value of the index for sector $j = 1, \dots, J = 10$ at time $i = 1, \dots, T$. Note that large positive values of $x_{i,j}$ indicate a large drop in the price index associated with sector j , and thus for risk management purposes we are interested in *large* values of $x_{i,j}$. Hence, for a given threshold u , we focus our attention on the collections of times $\{t_{j,k} : k = 1, \dots, n_j, j = 1, \dots, J\}$, where $t_{j,k}$ is the date associated with the appearance of the k -th negative log return in sector j that is larger than u .

For each sector j , we regard the collection of times $\{t_{j,k} : k = 1, \dots, n_j\}$ at which exceedances occur as a realization from a point process $N_j(t)$ defined on $[0, T]$. In turn, each $N_j(t)$ is constructed as the superposition of two independent, non-homogeneous Poisson processes. The first such process accounts for systematic risk and has a cumulative intensity function Λ_0^* that

is common to all sectors, while the second is associated with the idiosyncratic risk and has a cumulative intensity function Λ_j^* that is specific to each sector. Because of properties of superpositions of Poisson processes, this assumption implies that each $N_j(t)$ is also a non-homogeneous Poisson process with cumulative intensity $\Lambda_j(t) = \Lambda_0^*(t) + \Lambda_j^*(t)$, and intensity function $\lambda_j(t) = \lambda_0^*(t) + \lambda_j^*(t)$, where λ_0^* and λ_j^* are the Poisson process intensities associated with Λ_0^* and Λ_j^* , respectively.

The modeling approach for the Λ_j^* builds from the direct connection of a non-homogeneous Poisson process cumulative intensity/intensity function with a distribution/density function. Specifically, for $j = 0, 1, \dots, J$, we can write $\Lambda_j^*(t) = \gamma_j^* F_j^*(t)$, where $\gamma_j^* \equiv \Lambda_j^*(T) = \int_0^T \lambda_j^*(t) dt$ ($< \infty$) is the rate parameter controlling the total number of exceedances, and $F_j^*(t) = \Lambda_j^*(t)/\Lambda_j^*(T)$ is a distribution function on $[0, T]$ that controls how the exceedances are distributed over time. Hence, the sector-specific cumulative intensity function Λ_j can be written as

$$\Lambda_j(t) = \gamma_j F_j(t) = \{\gamma_0^* + \gamma_j^*\} \left\{ \frac{\gamma_0^*}{\gamma_0^* + \gamma_j^*} F_0^*(t) + \frac{\gamma_j^*}{\gamma_0^* + \gamma_j^*} F_j^*(t) \right\}.$$

Our construction implies that the sector-specific exceedance rate, γ_j , is simply the sum of the systematic and idiosyncratic rates, while the sector-specific distribution function, F_j , can be written as a mixture of the systematic and idiosyncratic distribution functions. The corresponding weight, $\epsilon_j = \gamma_0^*/(\gamma_0^* + \gamma_j^*)$, represents the proportion of exceedances in sector j that are associated with the systematic component. In addition, note that values of ϵ_j close to 1 (which are associated with $\gamma_0^* \gg \gamma_j^*$) imply a stronger association in the pattern of extreme losses.

Because each $N_j(t)$ follows a Poisson process, the probability that at most r exceedances will be observed in sector j during time period $[t_0, t_0 + \Delta]$ is

$$\sum_{i=0}^r \frac{\{\Upsilon_j(t_0, \Delta)\}^i \exp\{-\Upsilon_j(t_0, \Delta)\}}{i!},$$

where $\Upsilon_j(t_0, \Delta) = \Lambda_j(t_0 + \Delta) - \Lambda_j(t_0)$. These exceedance probabilities are easier to interpret than the intensity functions through which the model is defined. For example, the probability

that no exceedances are observed in sector j between time points t_0 and $t_0 + \Delta$ is given by

$$\begin{aligned} \exp \{ - [\Lambda_j(t_0 + \Delta) - \Lambda_j(t_0)] \} = \\ \exp \{ - [\Lambda_0^*(t_0 + \Delta) - \Lambda_0^*(t_0)] \} \times \exp \{ - [\Lambda_j^*(t_0 + \Delta) - \Lambda_j^*(t_0)] \}, \end{aligned}$$

where the first term in the right-hand side expression corresponds to the probability of no exceedance due to the systematic component and the second term corresponds to the probability of no exceedance due to the idiosyncratic component. In this sense, our model implies a multiplicative risk structure.

2.1 Modeling the intensity functions

In order to generate a flexible model that can capture changes in the pattern of extreme returns over time, we model the densities f_0^* and f_1^*, \dots, f_J^* associated with the systematic and idiosyncratic distribution functions F_0^* and F_1^*, \dots, F_J^* , using Dirichlet process mixtures. The proposed method builds on earlier work (Kottas and Behseta, 2010; Taddy and Kottas, 2012) on Dirichlet process mixture modeling for a single temporal Poisson process density. In particular, we let

$$f_j^*(t) \equiv f^*(t \mid G_j^*, \tau) = \int \psi(t \mid \mu, \tau) dG_j^*(\mu), \quad j = 0, 1, \dots, J,$$

where $\psi(t \mid \mu, \tau)$ is a kernel density on $[0, T]$ indexed by parameters μ and τ . The random mixing distribution G_j^* is assigned a Dirichlet process prior (Ferguson, 1973) with precision parameter α_j and baseline (centering) distribution H , which is common to all G_j^* . Thus, using the stick-breaking constructive definition of the Dirichlet process (Sethuraman, 1994),

$$G_j^*(\cdot) = \sum_{l=1}^{\infty} \left\{ v_{j,l} \prod_{s < l} (1 - v_{j,s}) \right\} \delta_{\tilde{\mu}_{j,l}}(\cdot), \quad j = 0, 1, \dots, J,$$

where $\delta_a(\cdot)$ denotes the degenerate measure at a , the atoms $\{\tilde{\mu}_{j,1}, \tilde{\mu}_{j,2}, \dots\}$ form a sequence of random variables independent and identically distributed according to H , and $\{v_{j,1}, v_{j,2}, \dots\}$ is another sequence of independent and identically distributed random variables according to a $\text{Beta}(1, \alpha_j)$ distribution.

Because in our application the support for the point process is a compact set, a natural choice for the kernel $\psi(t \mid \mu, \tau)$ is the rescaled beta density,

$$\frac{1}{T} \frac{\Gamma(\tau)}{\Gamma(\mu\tau/T) \Gamma(\{1 - \mu/T\}\tau)} \left(\frac{t}{T}\right)^{\mu\tau/T-1} \left(1 - \frac{t}{T}\right)^{\{1 - \mu/T\}\tau-1} \mathbb{I}_{[0,T]}(t) \quad (1)$$

where $\mu \in [0, T]$ is a location parameter (the mean of the kernel distribution), $\tau \in (0, \infty)$ can be interpreted as a scale parameter, and $\mathbb{I}_{\Omega}(\cdot)$ denotes the indicator function on the set Ω . Because Dirichlet process mixtures allow for a countable number of mixture components, the model is dense on the space of absolutely continuous distributions on $[0, T]$ as long as the baseline distribution H and the prior on the scale parameter τ are selected to provide full support on the domain of the parameters of the rescaled beta kernel (see, e.g., Diaconis and Ylvisaker, 1985). The precision parameter α_j controls the relative weight of the mixture components, with smaller values of α_j favoring mixtures with a small number of effective components. On the other hand, the baseline distribution H controls the location of the mixture components.

Besides a prior on the densities $f_0^*, f_1^*, \dots, f_J^*$, full prior specification for the intensity functions $\lambda_1(t), \dots, \lambda_J(t)$ requires priors for the rate parameters $\gamma_0^*, \gamma_1^*, \dots, \gamma_J^*$. In the case of the rate associated with the common component, γ_0^* , a natural choice is a gamma distribution with shape parameter $a_{\gamma_0^*}$ and rate parameter $b_{\gamma_0^*}$. For the idiosyncratic component we use a more general, zero-inflated gamma prior with density,

$$p(\gamma_j^* \mid \pi) = (1 - \pi)\delta_0(\gamma_j^*) + \pi \text{Gam}(\gamma_j^* \mid a_{\gamma_j^*}, b_{\gamma_j^*}), \quad j = 1, \dots, J$$

where $\text{Gam}(a, b)$ denotes a gamma distribution with mean a/b . Note that the case $\gamma_j^* = 0$ corresponds to $\epsilon_j = 1$, i.e., all exceedances in sector j are driven by systematic risks. Hence, this zero-inflated prior allows us to formally test for the presence of idiosyncratic risks. In the sequel we refer to this test as the *idiosyncracy test*.

Finally, we note that the representation of the sector-specific distribution function F_j as a mixture of a systematic and an idiosyncratic component is reminiscent of the models for dependent random distributions discussed in Müller et al. (2004) and Kolossiatis et al. (2013). In spite of this connection, the motivation for our modeling approach is different. Indeed, while

the motivation in the earlier work is to construct dependent random distributions, our model structure follows from the assumption that the point process associated with the presence of extreme returns can be constructed as a superposition of Poisson processes.

2.2 Hierarchical priors

The model is completed by adding hyperpriors to model parameters. For the baseline distribution H , we note that the parameter μ in Equation (1) is constrained to the $[0, T]$ interval. Hence, a natural choice for H is another rescaled beta distribution with density $h(\mu) \propto (\mu/T)^{a\mu-1} \{1 - (\mu/T)\}^{b\mu-1} \mathbb{I}_{[0, T]}(\mu)$. Although the full Bayesian model can be extended to incorporate a hyperprior for one or both of the parameters of H , this was not necessary for the application to the S&P500 data. In fact, we applied the model using the default choice of a uniform distribution for H ; sensitivity of posterior inference results to this choice, as well as to the other hyperprior choices, is discussed in Section 4.2.

The remaining priors are selected for computational convenience. In particular, a conditionally conjugate $\text{Gam}(a_\alpha, b_\alpha)$ prior is assigned to the precision parameters α_j , $j = 0, 1, \dots, J$. Similarly, the conditionally conjugate prior for the mixing probability π is given by a $\text{Beta}(a_\pi, b_\pi)$ distribution. Finally, the scale parameter τ is assigned an inverse gamma distribution with shape a_τ and scale b_τ (denoted by $\text{IGam}(a_\tau, b_\tau)$). Details on informative hyperparameter elicitation are discussed in Section 3.2

3 Computation and prior elicitation

3.1 Markov chain Monte Carlo algorithm

Our modeling approach is based on the decomposition of the Poisson process intensity into the total intensity, $\gamma_j = \gamma_0^* + \gamma_j^*$, and the Poisson process density, $f_j(t) = \epsilon_j f_0^*(t) + (1 - \epsilon_j) f_j^*(t)$, for $t \in [0, T]$. Under this decomposition, the likelihood function associated with the non-homogeneous Poisson process giving rise to the exceedances in sector j is $\gamma_j^{n_j} \exp\{-\gamma_j\} \prod_{k=1}^{n_j} f_j(t_{j,k})$. Hence,

the joint posterior distribution for our model can be written as

$$\begin{aligned}
p(\{\gamma_j^*\}, \{v_{j,l}\}, \{\tilde{\mu}_{j,l}\}, \tau, \{\alpha_j\}, \pi \mid \text{data}) &\propto \prod_{j=1}^J (\gamma_0^* + \gamma_j^*)^{n_j} \exp\{-(\gamma_0^* + \gamma_j^*)\} \\
&\times \prod_{j=1}^J \prod_{k=1}^{n_j} \left(\frac{\gamma_0^*}{\gamma_0^* + \gamma_j^*} \sum_{l=1}^{\infty} \left\{ v_{0,l} \prod_{s<l} (1 - v_{0,s}) \right\} \psi(t_{j,k} \mid \tilde{\mu}_{0,l}, \tau) + \right. \\
&\quad \left. \frac{\gamma_j^*}{\gamma_0^* + \gamma_j^*} \sum_{l=1}^{\infty} \left\{ v_{j,l} \prod_{s<l} (1 - v_{j,s}) \right\} \psi(t_{j,k} \mid \tilde{\mu}_{j,l}, \tau) \right) \\
&\times p(\pi) p(\tau) p(\gamma_0^*) \prod_{j=1}^J p(\gamma_j^* \mid \pi) \prod_{j=0}^J \prod_{l=1}^{\infty} h(\tilde{\mu}_{j,l}) \text{Beta}(v_{j,l} \mid 1, \alpha_j) \prod_{j=0}^J p(\alpha_j) \quad (2)
\end{aligned}$$

where the priors $p(\pi)$, $p(\tau)$, $p(\gamma_0^*)$, $p(\gamma_j^* \mid \pi)$, for $j = 1, \dots, J$, and $p(\alpha_j)$, for $j = 0, 1, \dots, J$, are given in Section 2.

Since the posterior distribution is computationally intractable, we resort to a Markov chain Monte Carlo (MCMC) algorithm (Robert and Casella, 2005) for simulation-based inference. To sample from the posterior distribution associated with the nonparametric component of the model, we resort to a blocked Gibbs sampler (Ishwaran and James, 2001). Hence, for computational purposes, we replace the mixing distributions $G_0^*, G_1^*, \dots, G_J^*$ with finite-dimensional approximations:

$$G_j^N(\cdot) = \sum_{l=1}^N \left\{ u_{j,l} \prod_{s<l} (1 - u_{j,s}) \right\} \delta_{\mu_{j,l}^*}(\cdot), \quad j = 0, 1, \dots, J.$$

where, as before, the $\mu_{j,l}^*$ are drawn independently from H and the $u_{j,l}$ are independently $\text{Beta}(1, \alpha_j)$ distributed, for $l = 1, \dots, N - 1$, but $u_{j,N} = 1$ to ensure that the weights sum to 1. The truncation level N can be chosen to any desired level of accuracy, using distributional results for partial sums of the Dirichlet process stick-breaking weights (see, e.g., Ishwaran and Zarepour, 2000); $N = 60$ was used for the results reported in Section 4.

Furthermore, we expand the model by introducing for each observation $t_{j,k}$ a pair of latent configuration variables, $(r_{j,k}, L_{j,k})$, where $r_{j,k} \in \{0, j\}$ is an indicator for the systematic or idiosyncratic component, whereas $L_{j,k} \in \{1, \dots, N\}$ identifies the respective mixture component under the truncation approximation for G_0^N or G_j^N . More specifically, independently for $j =$

$1, \dots, J$ and $k = 1, \dots, n_j$, $\Pr(r_{j,k} = 0 \mid \gamma_0^*, \gamma_j^*) = 1 - \Pr(r_{j,k} = j \mid \gamma_0^*, \gamma_j^*) = \gamma_0^*/(\gamma_0^* + \gamma_j^*)$, and, for $l = 1, \dots, N$, $\Pr(L_{j,k} = l \mid r_{j,k} = 0, G_0^N) = u_{0,l} \prod_{s < l} (1 - u_{0,s})$ and $\Pr(L_{j,k} = l \mid r_{j,k} = j, G_j^N) = u_{j,l} \prod_{s < l} (1 - u_{j,s})$. In addition to the aforementioned indicator variables associated with the mixture representation of the intensity functions, we introduce a set of binary indicators, ξ_1, \dots, ξ_J , such that $\Pr(\xi_j = 0 \mid \pi) = 1 - \pi$. Inferences on $\Pr(\xi_j = 0 \mid \text{data})$ provide an operational mechanism to implement the idiosyncrasy test discussed at the end of Section 2.1. Details of the posterior simulation algorithm are provided in the Appendix.

3.2 Hyperparameter elicitation

The model hyperparameters can be elicited using historical and/or expert information that is typically available for most liquid financial markets. We recommend that this elicitation process be complemented with a careful sensitivity analysis over a reasonable range of prior beliefs.

Consider first the parameters $\gamma_0^*, \gamma_1^*, \dots, \gamma_{10}^*$, which control the total number of exceedances observed in each sector and the relative distribution of these exceedances between the systematic and idiosyncratic component of the model. Because of their role in the model, we can elicit a value for the expected number of extreme returns in a given sector $j \geq 1$ (which corresponds to $\mathbf{E}\{\gamma_0^* + \gamma_j^*\}$) by assuming that returns in the sector are normally distributed, so that $\mathbf{E}\{\gamma_0^* + \gamma_j^*\} \approx T \times \Phi(\{-u - \zeta_j\}/\kappa_j)$, where Φ denotes the standard normal distribution function, and ζ_j and κ_j are rough estimates of the mean and standard deviation of returns for sector j . The values of ζ_j and κ_j can be obtained from historical data or expert knowledge. For simplicity, it can be assumed that ζ_j and κ_j are the same for every sector, leading to a model where sectors are exchangeable, but this is not required. Similarly, we can exploit the interpretation of $\gamma_0^*/(\gamma_0^* + \gamma_j^*)$ as the proportion of exceedances arising from the systematic component to elicit expert information about the most likely value of such rate, as well as a high probability range for its value. This same information can be used to provide informative priors for $1 - \pi$, the prior probability that the risk in a given sector is entirely driven by the systematic component.

Next, consider eliciting the hyperparameters associated with the densities f_0^* and f_1^*, \dots, f_J^* . A common feature of extreme returns in financial time series is that they tend to cluster over time (e.g., Mandelbrot, 1963). Hence, the priors for the precision parameters $\alpha_0, \alpha_1, \dots, \alpha_J$

(which, as mentioned in Section 2, control the number of mixture components) should favor a multimodal distribution. A rough value for the number of components, which can be used to select the prior mean of α_j (see, e.g., Escobar and West, 1995), can be elicited from a rough estimate of the frequency at which distress periods arise in sector j . Similarly, the value for the scale parameter τ can be elicited from prior information about the length of distress periods. Finally, in the absence of prior information about the time at which distress periods occur, we recommend that H be selected so that the prior mean for f_0^* and f_1^*, \dots, f_J^* is close to uniform.

4 Systematic and idiosyncratic risks in the U.S. economy: An analysis of the S&P500 index and its components

The data analyzed in this section corresponds to the negative daily log returns above $u = 2\%$ on each of the ten sectors that make up the S&P500 index, from the time period between January 1, 2000 and December 31, 2011. This results in sample sizes for extreme returns that range from 85 (for consumer staples) to 387 (for information technology). Prices for the individual indexes were obtained from Bloomberg financial services; the corresponding tickers are S5COND, S5CONS, S5ENRS, S5FINL, S5HLTH, S5INDU, S5MATR, S5INFT, S5TELS, S5UTIL. A 2% drop in the market has been historically used as a threshold for trading curbs on program trades (which involve a basket of stocks from the S&P500 index where there are at least 15 stocks or where the value of the basket is at least \$1 million). Although these particular trading curbs were abandoned in late 2007 because of their inefficacy in reducing market volatility, we believe that the 2% threshold is still a useful guideline to identify periods of distress in the market without excessively thinning the sample. We repeated our analysis for other values of the threshold u without significant qualitative changes in the conclusions.

All inferences reported here are based on 3,000 posterior samples obtained after a burn in period of 20,000 iterations and thinning of the original chain every 50 iterations. Convergence of the MCMC algorithm was monitored using trace plots as well as the R statistic discussed in Gelman and Rubin (1992). In particular, we ran four independent chains started from overdispersed initial values and compared between and within chain variability in the value of the log

likelihood function and in some of the hyperparameters in the model. No lack of convergence was evident from these diagnostics. The algorithm was implemented in C/C++ and total execution time was approximately 16 hours on a MacBook laptop with a 2 GHz Intel Core 2 Duo processor and 2GB of memory.

Following the approach described in Section 3.2, we chose the hyperparameters for γ_0^* and γ_j^* so that $a_{\gamma_0^*} = 7.32$, $b_{\gamma_0^*} = 0.06$, and $a_{\gamma_j^*} = 1.32$, $b_{\gamma_j^*} = 0.06$. These values were chosen on the basis of a 0 mean return with 18% annualized volatility for the S&P500, along with the prior beliefs that on average 85% of the observed exceedances, and with .99 probability at least 50% of them, arise from the systematic component of the model. These numbers reflect our prior belief that the systematic component of the risk will explain a majority of the exceedances observed in the data. This is further emphasized by our choice for the hyperparameters for π , which are selected as $a_\pi = 0.5$ and $b_\pi = 2$. This choice implies that $E\{\pi\} = 0.2$ and places a high probability on values of π close to zero. Finally, the prior for the scale parameter τ was selected so that $a_\tau = 5$ and $b_\tau = 2,400$ (which implies that $E\{\tau\} = 600$); for the Dirichlet process precision parameters, we chose $a_\alpha = 4$ and $b_\alpha = 1/3$ (which favors highly multimodal intensity functions a priori); and the parameters for the baseline distribution were set to $a_\mu = 1$ and $b_\mu = 1$. The posterior densities for all hyperparameters appear to be concentrated relative to the corresponding prior densities. In addition, we note that moderate changes in these assumptions (e.g., assuming that on average only 50% of the exceedances come from the systematic component, or a 2% positive annualized return for the S&P500 with a 25% annualized volatility, or a uniform prior on π) lead to essentially equivalent posterior inferences for the intensity functions. The only inferences that are somewhat affected by moderate changes in the prior hyperparameters are those associated with the idiosyncrasy test discussed in Section 2; for further details on prior sensitivity analysis see Section 4.2.

Estimates of the overall intensities $\lambda_1(t), \dots, \lambda_{10}(t)$ associated with each of the ten components of the S&P500 index can be seen in Figure 2. The last two panels also provide summaries of the prior distribution over intensities induced by the prior choices discussed above. By comparing some of those estimates to the raw data presented in Figure 1 it becomes clear that the model faithfully reproduces the main features of the data. Furthermore, the uncertainty asso-

ciated with these estimates is relatively low. We note that these intensities (or perhaps their ratios with respect to the intensity associated with the systematic component, $\frac{\lambda_1(t)}{\lambda_0(t)}, \dots, \frac{\lambda_{10}(t)}{\lambda_0(t)}$) can potentially be used as a time-dependent index of idiosyncratic risk on each of the sectors of the U.S. economy tracked by the S&P500 index.

[Figure 2 about here.]

Next, Figures 3 and 4 show estimates of the densities f_0^* and f_1^*, \dots, f_{10}^* associated with the systematic and idiosyncratic risk intensities. Figure 5 presents the posterior densities for $\epsilon_1, \dots, \epsilon_{10}$, the (average) proportion of the risk attributable to the systematic component in each of the ten sectors. Note that in half the sectors (consumer discretionary, consumer staples, health care, industrials, and utilities) the proportion of extremes associated with the systematic component is at least 80%, while for the rest (energy, financials, information technology, telecommunications and materials) the proportion is between 40% and 80%. In addition, note that the density for the systematic risk shows peaks that coincide, or shortly follow, important stock market events. On the other hand, the behavior of the idiosyncratic risk varies drastically with the economic sector and, in most cases, can be explained by factors that are sector-specific. For example, the energy and utilities sectors present increases in idiosyncratic risk during 2005, a period that corresponded to sharp increases in oil prices but that was otherwise relatively calm. On the other hand, the idiosyncratic risk associated with the financial sector increases dramatically after the summer of 2007. An oversized idiosyncratic risk for this sector after 2007 is clearly reasonable as financials were the main driver of the recent crisis. Similarly, the idiosyncratic risks associated with the information technology and telecommunication services sectors are particularly elevated between 2000 and 2002, a period that included the bursting of the so-called dot-com bubble. Finally, note that the idiosyncratic risk associated with consumer staples is almost negligible over the whole period under study, with our idiosyncrasy test suggesting that there is moderate evidence for the absence of idiosyncratic risk in this sector of the S&P500 index. This is reasonable, as the consumer staples sector includes companies that produce and trade basic necessities whose consumption might be affected by general economic conditions but is otherwise relatively stable.

[Figure 3 about here.]

[Figure 4 about here.]

[Figure 5 about here.]

As we discussed in Section 2, we can alternatively quantify the level of risk through the probability of observing at least one exceedance during a given period of time. Figure 6 shows for four different sectors the posterior distributions for the odds ratios of the probability of at least one exceedance in the month starting two weeks after the bankruptcy of Lehman Brothers versus the probability of at least one exceedance in the month ending two weeks before the bankruptcy. Note that all sectors show an increase in risk after Lehman Brothers bankruptcy. However the increase is lower for financials than it is for the other sectors (the estimated posterior probabilities are 1.000, 0.913 and 0.977 for consumer staples, energy and information technology, respectively). Indeed, note that systematic risk increases after the bankruptcy of Lehman Brothers but the idiosyncratic risk associated with financials actually decreases (as does the one for energy, although to a lesser degree), while the idiosyncratic risks associated with information technology and consumer staples increased. The increase in risk in the information technology and consumer staples sectors could be explained by the fact that one of the main effects of Lehman's bankruptcy was a collapse in the short term corporate debt market. Hence, although the bankruptcy of Lehman Brothers actually reduced the uncertainty in the financial sector of the economy, it caused real damage to companies in other sectors that are extremely dependent on short term debt. Note that companies that are part of the S&P500 energy sector are typically not reliant in short term funding, hence the limited impact of Lehman's bankruptcy in their idiosyncratic risk.

[Figure 6 about here.]

4.1 Model validation

The model was validated using two different approaches. First, an out-of-sample cross validation exercise was conducted, with cross-validation datasets being constructed by randomly selecting

20% of the observations from each sector to be used as held-out data. The remaining 80% of the data was used to fit our nonparametric model and generate nominal 90% highest posterior density (HPD) intervals for new exceedances. The true coverage of these HPD intervals was then evaluated on the held-out data. Figure 7 presents examples of cross-validation samples and the corresponding densities for two different sectors.

[Figure 7 about here.]

We repeated the process described above for 10 different cross-validation datasets, with the results being presented in Figure 8. As expected, there is variability in the coverage rates depending on the sector and the specific cross-validation dataset. However, the results suggest that for the most part the real coverage rates are in line with the nominal coverage, which suggest that the model does not under or over fit.

[Figure 8 about here.]

In addition to the cross validation exercise described above, in-sample goodness of fit was investigated using quantile-quantile plots for the posterior distribution of inter-event times. More specifically, we use the time-rescaling theorem (see, e.g., Daley and Vere-Jones, 2003), according to which if $\{t_{j,k} : k = 1, \dots, n_j\}$ is a realization from a non-homogeneous Poisson process with cumulative intensity $\Lambda_j(t)$, then the transformed point pattern $\{\Lambda_j(t_{j,k}) : k = 1, \dots, n_j\}$ is a realization from a homogeneous Poisson process with unit intensity. Hence, if the Poisson process assumption is correct, we would expect the transformed inter-arrival times $z_{j,1}, \dots, z_{j,n_j}$ defined as $z_{j,k} = 1 - \exp\{-[\Lambda_j(t_{j,k}; G_j, \tau) - \Lambda_j(t_{j,k-1}; G_j, \tau)]\}$ (with the convention $t_{j,0} = 0$) to be uniformly distributed on the unit interval. A variant of this check for extreme value models was originally proposed by Hill (1975) in the context of threshold selection.

Figure 9 presents quantile-quantile plots of the posterior means of these transformed inter-arrival times, $E\{z_{j,k} \mid \text{data}\}$, for $k = 1, \dots, n_j$, against the quantiles of a uniform distribution for each of the ten S&P500 sectors. For the most part the fit appears acceptable, although there is some evidence of poor fit for a couple of sectors. In particular, note that for consumer staples our model tends to systematically predict shorter inter-arrival periods than those that would

be expected under the Poisson model. Similar, albeit less dramatic, biases can also be seen for information technology.

[Figure 9 about here.]

4.2 Sensitivity analysis

We carried out a comprehensive sensitivity analysis to assess the effect of prior distributions on posterior inferences. First, we considered three alternative sets of hyperparameters for the mixture kernel scale parameter τ corresponding to $\text{IGam}(5, 4000)$, $\text{IGam}(10, 7200)$, and $\text{IGam}(2, 500)$ priors. The hyperparameters were selected to represent a range of situations where the prior mean is both larger and smaller than the one used for our previous analysis, as well as different levels of concentration. Posterior inferences were mostly unaffected under any of these scenarios.

Next, we considered four alternative prior specifications for the precision parameters $\alpha_0, \dots, \alpha_{10}$, specifically, a $\text{Gam}(4, 1/3)$, a $\text{Gam}(10, 2)$, a $\text{Gam}(2, 0.4)$, and a $\text{Gam}(3, 3)$ prior. Inferences for the intensity function were mostly unchanged under these prior distributions. However, inferences for individual hyperparameters were somewhat affected. In particular, smaller values for $E\{\alpha_j\}$ naturally lead to somewhat smaller posterior means for the α_j , but also to larger posterior values for τ and an increase in the posterior mean for some of the ϵ_j (for example, for consumer staples we have $\Pr(\epsilon_2 = 1 \mid \text{data}) = 0.71$ under the $\text{Gam}(3, 3)$ prior). On the other hand, changes in the prior dispersion of α_j had no discernible effect on the individual posterior distributions of the hyperparameters, as long as the prior mean was kept constant.

To assess the effect of the baseline distribution H on posterior inferences, we considered an alternative rescaled beta distribution with $a_\mu = b_\mu = 3$. This prior tends to favor the localization of distress periods towards the middle of the time series. This alternative baseline distribution leads to somewhat smoother estimates for the density functions f_1, \dots, f_{10} , and to less multimodal estimates for the idiosyncratic intensities f_1^*, \dots, f_{10}^* . In addition, the posterior means for the α_j and for τ tend to be slightly smaller. However, the overall structure of the results is unchanged.

We also considered an alternative specification for the priors on γ_0^* and $\gamma_1^*, \dots, \gamma_{10}^*$ where we assume that the number of extreme returns in each sector is consistent with a 2% positive annu-

alized return and a 25% annualized volatility for the S&P500, while only 50% of the exceedances come from the systematic component, and $\Pr(0.2 < \gamma_0^*(\gamma_0^* + \gamma_j^*)^{-1} < 0.8) = 0.99$. This leads to $a_{\gamma_0^*} = 7.65$, $b_{\gamma_0^*} = 0.65$, and $a_{\gamma_j^*} = 7.65$, $b_{\gamma_j^*} = 0.65$ for $j \geq 1$. As before, these new priors have very little impact on the inference for the intensity functions. However, the inferences on the weights $\epsilon_1, \dots, \epsilon_{10}$ are significantly affected. In particular, the idiosyncrasy test now provides very strong evidence that shocks in consumer staples are driven exclusively by the systematic component ($\Pr(\epsilon_2 = 1 \mid \text{data}) \approx 1$), while for the other sectors we have very strong evidence for the presence of idiosyncratic components ($\Pr(\epsilon_j = 1 \mid \text{data}) = 0$ for all $j = 1, 3, 4, \dots, 10$).

Finally, we investigated the effect on posterior inferences of alternative prior distributions on π . In addition to the original $\text{Beta}(0.5, 2)$ prior (which favors the hypotheses that most the exceedances are generated by the systematic component of the model), we considered a uniform and a $\text{Beta}(0.5, 0.5)$ prior for π . While the $\text{Beta}(0.5, 0.5)$ had a negligible effect on posterior inferences, the use of a uniform prior led to an increase in the posterior mean for some of the ϵ_j (for example, $\Pr(\epsilon_2 = 1 \mid \text{data}) = 0.76$).

5 Discussion

We have developed a novel modeling approach for simultaneous risk assessment in multiple sectors of the U.S. economy that focuses on modeling the dependence in the appearance of large losses in the different sectors constituting the S&P500 index. Some of the advantages of the proposed methodology include its nonparametric nature, its focus on extreme returns rather than average returns, and the interpretability of the components of the model. These features lead to robust indexes and hypothesis tests for the presence of idiosyncratic risks in specific sectors of the economy.

Although our application used data from equity markets to understand the interdependence across different sectors of the economy, the model described here also has potential applications in debt markets. In particular, reduced-form credit risk models (e.g., Lando, 1998) are also Cox process models for debt default. Therefore, our model could be extended to estimate the probability of default for firms within and across multiple economic sectors.

Appendix. Details of the Markov chain Monte Carlo algorithm

Here, we provide the details for simulation from the posterior distribution of the model given in Equation (2). Posterior simulation is based on the blocked Gibbs sampler, including Metropolis-Hastings steps for some of the parameters. In particular, the MCMC algorithm iteratively updates model parameters through the following steps:

1. Each pair $(r_{j,k}, L_{j,k})$, for $j = 1, \dots, J$ and $k = 1, \dots, n_j$, is updated jointly from a multinomial distribution with

$$\Pr(r_{j,k} = 0, L_{j,k} = l \mid \dots, \text{data}) \propto \gamma_0^* \left[u_{0,l} \prod_{s < l} \{1 - u_{0,s}\} \right] \psi(t_{j,k} \mid \mu_{0,l}^*, \tau)$$

and

$$\Pr(r_{j,k} = j, L_{j,k} = l \mid \dots, \text{data}) \propto \gamma_j^* \left[u_{j,l} \prod_{s < l} \{1 - u_{j,s}\} \right] \psi(t_{j,k} \mid \mu_{j,l}^*, \tau),$$

for $l = 1, \dots, N$.

2. The update for the stick-breaking ratios $\{u_{j,l}\}$ follows the standard recipe described in Ishwaran and James (2001). In particular, the posterior full conditional for each $u_{j,l}$ is a Beta $\left(1 + M_{j,l}, \alpha_j + \sum_{r=l+1}^N M_{j,r}\right)$ distribution, where $M_{j,l} = \sum_{s=1}^J \sum_{i=1}^{n_k} \mathbb{I}_{(r_{s,i}=j, L_{s,i}=l)}$.
3. The full conditional distribution for the atoms $\mu_{j,l}^*$ with $j = 0, 1, \dots, J$ and $l = 1, \dots, N$ is given by

$$p(\mu_{j,l}^* \mid \dots, \text{data}) \propto h(\mu_{j,l}^*) \prod_{\{(s,i): r_{s,i}=j, L_{s,i}=l\}} \psi(t_{s,i} \mid \mu_{j,l}^*, \tau).$$

Hence, if $\{(s, i) : r_{s,i} = j, L_{s,i} = l\} = \emptyset$, $\mu_{j,l}^*$ can be updated by sampling from the baseline distribution H . Otherwise, since this full conditional does not correspond to any known distribution, we update $\mu_{j,l}^*$ using a random walk Metropolis-Hasting step with symmetric logit-normal proposal distribution where new values $\mu_{j,l}^{*(p)}$ are generated according to $\text{logit} \left\{ \mu_{j,l}^{*(p)} \right\} \mid \mu_{j,l}^{*(c)} \sim \text{Normal} \left(\text{logit} \left\{ \mu_{j,l}^{*(c)} \right\}, \kappa_\mu^2 \right)$, where $\mu_{j,l}^{*(c)}$ represents the current value of the chain, and κ_μ^2 is a tuning parameter chosen such that the average acceptance rate is between 30% and 40%.

4. Since the prior for γ_0^* is conditionally conjugate, we update this parameter by sampling from $\gamma_0^* \mid \dots, \text{data} \sim \text{gamma} \left(a_{\gamma_0^*} + \sum_{j=1}^J m_j, b_{\gamma_0^*} + J \right)$, where $m_j = \sum_{k=1}^{n_j} \mathbb{I}_{(r_{j,k}=0)}$ denotes the number of observed exceedances in sector j that are associated with the systematic component.
5. For $j = 1, \dots, J$, the pairs (ξ_j, γ_j^*) are conditionally independent from each other and can be updated by first updating ξ_j so that $\xi_j = 1$ if $m_j < n_j$ or $\xi_j \mid \dots, \text{data} \sim \text{Bernoulli} \left(\frac{\pi \left\{ b_{\gamma_j^*} / (1 + b_{\gamma_j^*}) \right\}^{a_{\gamma_j^*}}}{1 - \pi + \pi \left\{ b_{\gamma_j^*} / (1 + b_{\gamma_j^*}) \right\}^{a_{\gamma_j^*}}} \right)$ if $m_j = n_j$. Once ξ_j has been updated, γ_j^* is updated by setting $\gamma_j^* = 0$ if $\xi_j = 0$, or by sampling γ_j^* from a gamma distribution with shape parameter $a_{\gamma_j^*} + n_j - m_j$ and rate parameter $1 + b_{\gamma_j^*}$ if $\xi_j = 1$.
6. The full conditional posterior for τ is

$$p(\tau \mid \dots, \text{data}) \propto p(\tau) \prod_{j=1}^J \prod_{k=1}^{n_j} \psi(t_{j,k} \mid \mu_{r_{j,k}^*, L_{j,k}}, \tau).$$

Since no direct sampler is available from this distribution, we update τ using a random walk Metropolis-Hasting step with symmetric log-normal proposal, $\log(\tau^{(p)}) \mid \tau^{(c)} \sim \text{Normal}(\log(\tau^{(c)}), \kappa_\tau^2)$, where κ_τ^2 is again a tuning parameter.

7. Because the prior on π is conditionally conjugate, the full conditional for π is given by a $\text{Beta} \left(a_\pi + J - \sum_{j=1}^J \xi_j, b_\pi + \sum_{j=1}^J \xi_j \right)$ distribution.
8. The precision parameters $\alpha_0, \alpha_1, \dots, \alpha_J$ can be updated independently using the algorithm described in Escobar and West (1995).

References

- Allen, D., P. Gerrans, A. Singh, and R. Powell (2009). Quantile regression: Its application in investment analysis. *Journal of the Securities Institute of Australia* **1**, 1–12.
- Barnes, M. and W. Hughes (2002). A quantile regression analysis of the cross section of stock market returns. Technical report, Federal Reserve Bank of Boston.
- Chang, M., J. Hung, and C. Nieh (2011). Reexamination of capital asset pricing model (capm): An application of quantile regression. *African Journal of Business Management* **5**, 12684–12690.

- Cox, D. R. (1955). Some statistical methods connected with series of events. *Journal of the Royal Statistical Society, Series B* **17**, 129–164.
- Daley, D. J. and D. Vere-Jones (2003). *An Introduction to the Theory of Point Processes* (Second ed.). Springer, New York.
- Daniel, K. D., D. Hirshleifer, and A. Subrahmanyam (2001). Overconfidence, arbitrage, and equilibrium asset pricing. *The Journal of Finance* **56**, 921–965.
- Diaconis, P. and D. Ylvisaker (1985). Quantifying prior opinion. In J. Bernardo, M. DeGroot, D. Lindley, and A. Smith (Eds.), *Bayesian Statistics 2*, pp. 133–156. Amsterdam, North-Holland: Oxford University Press.
- Escobar, M. D. and M. West (1995). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association* **90**, 577–588.
- Fama, E. F. and K. R. French (1992). The cross-section of expected stock returns. *Journal of Finance* **47**, 427–465.
- Fama, E. F. and K. R. French (2004). The capital asset pricing model: Theory and evidence. *Journal of Economic Perspectives* **18**, 25–46.
- Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *The Annals of Statistics* **1**, 209–230.
- French, C. W. (2003). The Treynor capital asset pricing model. *Journal of Investment Management* **1**, 60–72.
- Gelman, A. and D. Rubin (1992). Inferences from iterative simulation using multiple sequences. *Statistical Science* **7**, 457–472.
- Hill, B. M. (1975). A simple general approach to inference about the tail of a distribution. *Annals of Statistics* **3**, 1163–1174.
- Ishwaran, H. and L. F. James (2001). Gibbs sampling methods for stick-breaking priors. *Journal of the American Statistical Association* **96**, 161–173.
- Ishwaran, H. and M. Zarepour (2000). Markov chain Monte Carlo in approximate Dirichlet and Beta two-parameter process hierarchical models. *Biometrika* **87**, 371–390.
- Kolossiatis, M., J. E. Griffin, and M. F. J. Steel (2013). On Bayesian nonparametric modelling of two correlated distributions. *Statistics and Computing* **23**, 1–15.
- Kottas, A. and S. Behseta (2010). Bayesian nonparametric modeling for comparison of single-neuron firing intensities. *Biometrics* **66**, 277–286.
- Lando, D. (1998). On Cox processes and credit risk securities. *Review of Derivatives Research* **2**, 99–120.

- Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business* **36**, 394–419.
- Müller, P., F. Quintana, and G. Rosner (2004). Hierarchical meta-analysis over related non-parametric Bayesian models. *Journal of Royal Statistical Society, Series B* **66**, 735–749.
- Robert, C. P. and G. Casella (2005). *Monte Carlo Statistical Methods* (Second Edition ed.). Springer, New York.
- Sethuraman, J. (1994). A constructive definition of Dirichlet priors. *Statistica Sinica* **4**, 639–650.
- Taddy, M. and A. Kottas (2012). Mixture modeling for marked Poisson processes. *Bayesian Analysis* **7**, 335–362.
- Treynor, J. L. (1961). Market value, time, and risk. Unpublished manuscript.
- Treynor, J. L. (1962). Toward a theory of market value of risky assets. Unpublished manuscript.




Figure 1: Negative log returns above 2% for four sectors of the S&P500 index (consumer staples, energy, financials, and information technology). Vertical dotted lines identify seven events of significance to the markets: the bursting of the .com bubble (03/10/2000), the 9/11 terrorist attacks (09/11/2001), the stock market downturn of 2002 (09/12/2002), the bursting of the Chinese bubble (02/27/2007), the bankruptcy of Lehman Brothers (09/16/2008), Dubai’s debt standstill (11/27/2009), and the beginning of the European sovereign debt crisis (08/27/2010).




Figure 2: Posterior mean of the overall intensity associated with the different components of the S&P500 index, along with posterior 95% pointwise credible intervals. The headers on each panel include the number of exceedances observed in each sector over the 12 year period under study. The last two figures in the bottom row present prior realizations for the intensity function (central panel) and the mean prior intensity function along with prior 95% pointwise credible intervals (right panel).




Figure 3: The left panel shows the posterior mean of the density associated with the systematic risk component of the S&P500 index, including posterior 95% pointwise credible intervals. The central panel shows realizations from the prior density function, while the right panel shows the mean prior density function and prior 95% pointwise credible intervals.




Figure 4: Posterior mean of the densities associated with the different components of the S&P500 index, along with posterior 95% pointwise credible intervals.

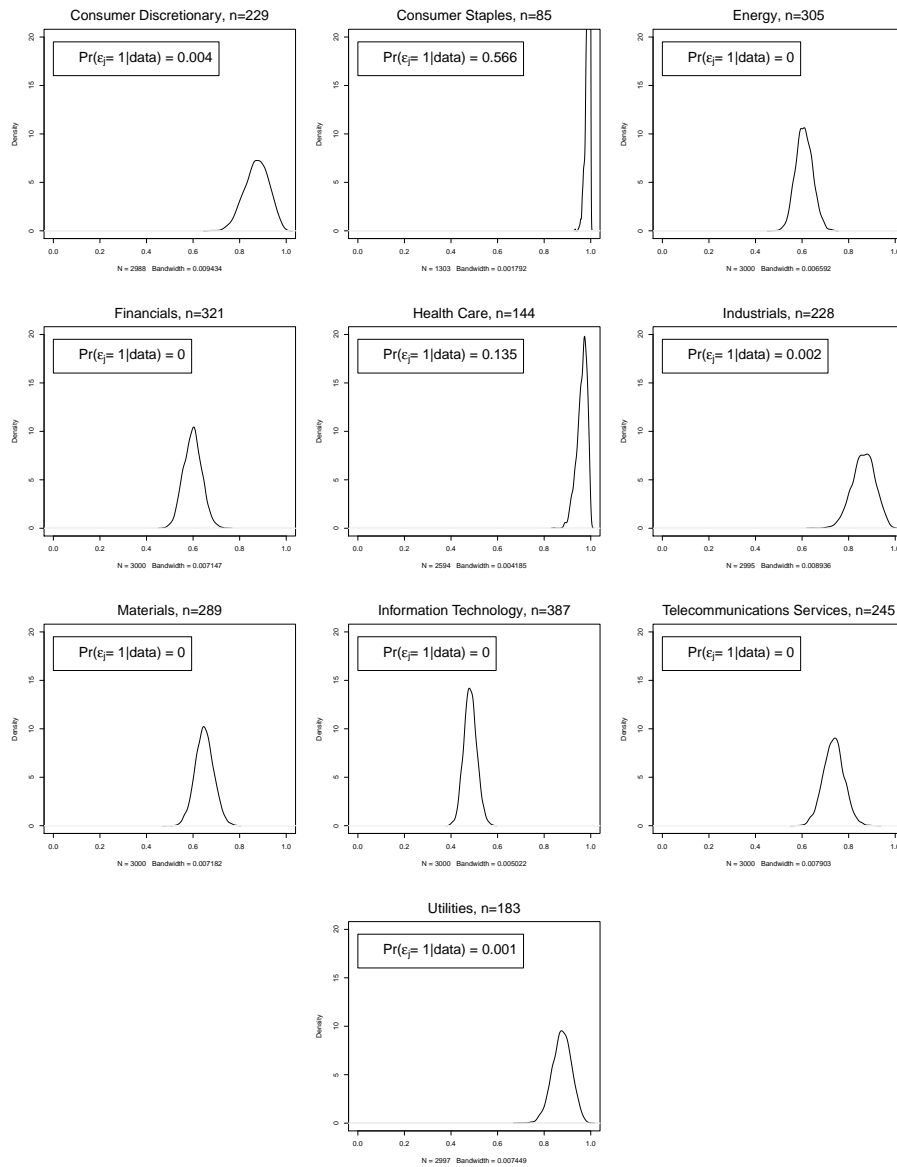


Figure 5: Posterior densities for the overall proportion of risk attributable to the systematic component on each of the ten components of the S&P500 index.

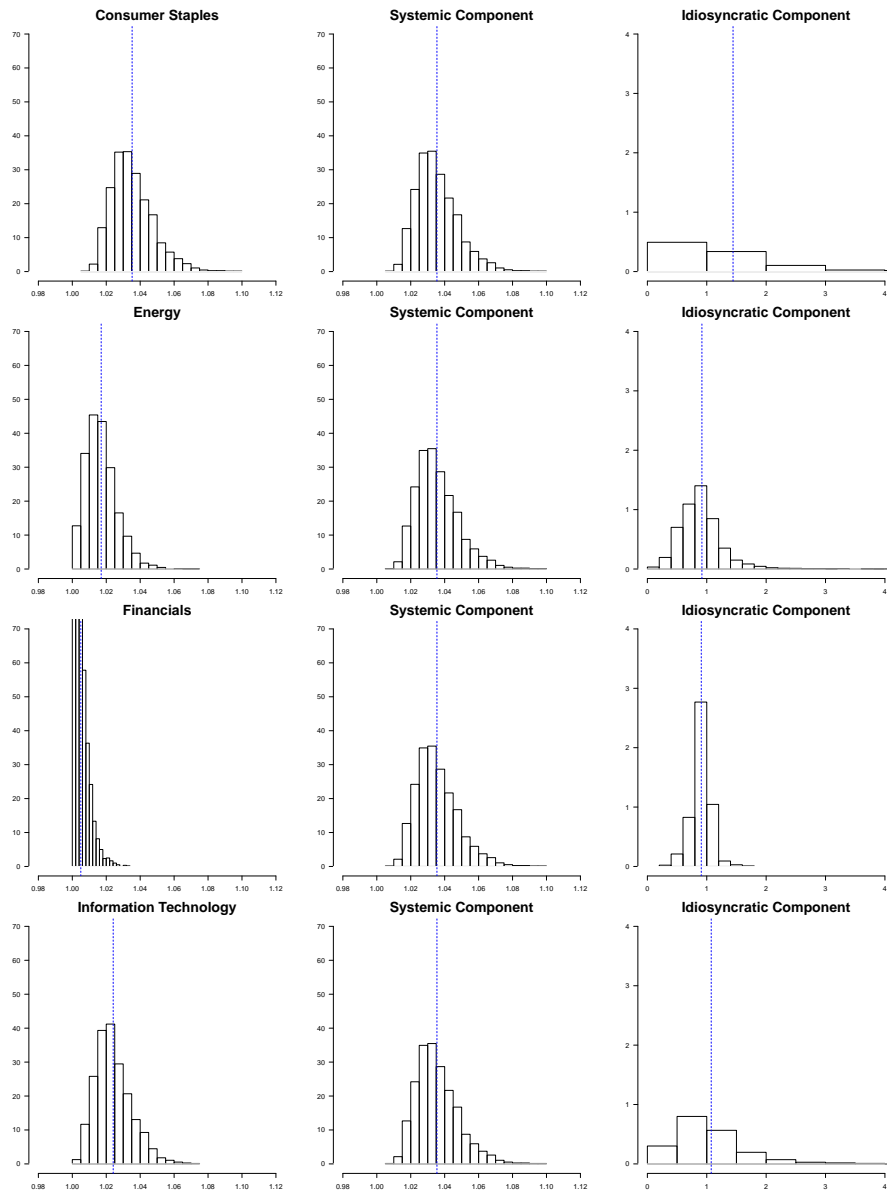


Figure 6: Posterior distributions for the odds ratio of the probability of at least one exceedance in the month starting two weeks after the bankruptcy of Lehman Brothers against the probability of at least one exceedance in the month ending two weeks before the bankruptcy for four different sectors. The vertical line corresponds to the mean of the posterior distribution.




Figure 7: Examples of cross-validation datasets (plotted on the horizontal axis) and the density estimates associated with them.




Figure 8: Results from the cross validation exercise to investigate the coverage rate of highest posterior density intervals associated with the nonparametric model.




Figure 9: Quantile-quantile plot of the posterior means for the transformed inter-arrival times against the quantiles of a uniform distribution (solid line) for each of the ten S&P500 sectors.