

Recovering Conditional Return Distributions by Regression: Estimation and Applications*

Fang Liu[†]
Job Market Paper

December 2014

Abstract

I propose a regression approach to recovering the return distribution of an individual asset conditional on the return of an aggregate index based on their marginal distributions. This approach relies on the identifying assumption that the conditional return distribution of the asset given the index return does not vary over time. I show how to empirically implement this approach using option price data. I then apply this approach to examine the cross-sectional equity risk premium associated with systematic disaster risk, to estimate the exposure of banks to systemic shocks, and to extend the Ross (*Journal of Finance*, 2014) recovery theorem to individual assets.

1 Introduction

The recent financial crisis has witnessed dramatic declines in the prices of most securities, which suggests strong return comovement between various assets. It is desirable to understand how the returns of different securities move along with each other. In this paper, I propose a regression approach to recovering the return distribution of an individual asset conditional on the return of an aggregate index based on their marginal distributions.

*I thank Jason Donaldson, Phil Dybvig, George Jiang, Ohad Kadan, Isaac Kleshchelski, Jeongmin Lee, Mark Loewenstein, Asaf Manela, Thomas Maurer, Giorgia Piacentino, Matthew Ringgenberg, Stephen Ross, Ngoc-Khanh Tran, Guofu Zhou, as well as seminar participants at the Institute of Financial Studies at Southwestern University of Finance and Economics, China, and Washington University in St. Louis for helpful comments and suggestions.

[†]Olin Business School, Washington University in St. Louis. E-mail: fliu23@wustl.edu.

I show how this approach can be implemented empirically using option prices. I then demonstrate the usefulness of this approach in testing the cross-sectional equity premium associated with systematic disaster risk, in estimating the exposure of banks to systemic shocks, and in an extension of the Ross (2014) recovery theorem.

The simplest and most widely used approach to describing the joint return behavior between two securities is to run a linear regression based on their historical returns. Indeed, this is what we are used to doing when estimating the CAPM beta by regressing excess returns of an asset on those of the market portfolio. This approach, however, has a number of drawbacks. First, it estimates the mean return of one security conditional on the return of the other, but it fails to capture high-moment properties. For example, Figure 1 plots the returns of two pairs of hypothetical securities, both of which predict the same conditional mean return of one security given the return of the other. However, the second-moment patterns of the two pairs are clearly distinct in the sense that the first pair has increasing correlation when the returns become lower, whereas the second pair has symmetric correlation over the entire region of returns. Second, running a linear regression between the two securities focuses on the linear relation only, neglecting other aspects of their joint behavior. To illustrate this point, Figure 2 depicts the returns of two hypothetical assets against the return of the market portfolio. The returns of both assets fit the same linear relation with the market return. Nevertheless, the nonlinear patterns show that asset 1 is more sensitive to the market disaster risk than asset 2 in the sense that the former tends to deliver lower returns when the market return becomes disastrously low. Third, estimation based on historical returns is backward-looking, which does not necessarily represent future return distributions. Finally, using historical returns makes it difficult to capture the effects of rare events, especially when the sample size is not large enough.

Alternative approaches used in the literature resolve some of the above issues. One such approach is the quantile regression, which predicts the conditional return quantiles of one security given the return of the other. (See Koenker (2005) for detailed discussions on quantile regression.) This approach generates the entire conditional distributions, thus capturing high-moment properties as well as nonlinear aspects in the joint return behavior. Nevertheless, given that the quantile regression is also implemented using historical returns,

it is backward-looking and does not adequately reflect the effects of rare events.

Another alternative relies on option prices. We learn from the results of Ross (1976) and Breeden and Litzenberger (1978) that one can estimate the risk-neutral probability distribution of security returns using prices of options written on the security under consideration. The advantage of using option prices is that it is forward-looking and accounts for rare events even if such events do not occur within sample. However, the risk-neutral distributions obtained from option prices typically differ from the physical distributions due to the adjustment for risk aversion. In addition, if we are interested in the joint return distribution of two securities, we would need options written on the joint values of these two securities. Given that most traded options are written on a single security, this method generally allows one to estimate the risk-neutral marginal return distribution of each single security, but not their joint distribution.

A question that follows naturally is whether we can recover the joint return distribution of two securities, assuming that the associated marginal distributions are known or can be estimated. This is indeed straightforward in the special cases in which the returns of the two securities are perfectly correlated or independent of each other. For more general cases, a well-known tool for this purpose is the copula, which can be used to map the marginal return distributions of multiple securities to their joint distribution. (See Nelsen (1999) for a general overview of the copula method.) However, a drawback of this approach is that it is parametric in the sense that it typically relies on specifying a particular class of copulas. When the copula class is misspecified, the accuracy of estimation might be affected.

In light of all the problems discussed above, it is desirable to have a better approach to evaluating the joint return behavior of two different securities. The term “better” includes the following aspects. First, it should capture all moment properties of the joint return distribution. Second, it should capture linear as well as nonlinear relations in the returns of the two securities. Third, it should reflect forward-looking information. Fourth, it should naturally account for rare events, whose ex-ante probabilities of occurrence are extremely small. Finally, it should not depend on any parametric assumptions on the return distributions of the securities.

I propose a novel approach of recovering the conditional return distribution of an individual asset given the return of an aggregate index from their marginal distributions. The

index return can be viewed as a factor that determines the state of the economy. Examples of the aggregate index include the market portfolio or a sector portfolio, etc. According to the total probability formula, the marginal return distributions of the two securities are linked to each other through the conditional return distribution of the asset given the index return. I assume that the conditional return distribution of the asset given any particular value of the index return remains fixed over time, meaning that the time variation in the return distribution of the asset is solely driven by that of the index. This allows me to estimate the time-invariant conditional return probabilities of the asset as the coefficients from a constrained linear regression of the marginal return distribution of the asset on that of the index over time. I show that under the standard OLS assumptions, the estimates from this constrained regression are consistent, i.e., they converge to the true conditional probabilities as the sample size becomes large enough.

I further assume that the variation in the index return is the only priced risk (systematic risk) such that any variation left in the asset return is idiosyncratic and does not get priced. This implies that the conditional return distribution of the asset given the index return is the same under the physical and the risk-neutral probability measures. Since risk-neutral marginal distributions of security returns can be extracted from option prices, my approach can be implemented under the risk-neutral measure using option pricing data. The resulting conditional return distribution of the asset estimated this way coincides with the physical conditional distribution.

The advantages of my approach include the following. First, it generates the entire conditional return distribution of the asset given the index return, thus capturing all moment properties and potential nonlinearity in their joint behavior. In addition, since this approach can be implemented using option prices, it is forward-looking and accounts for the likelihood of rare events perceived by investors even if such events do not truly occur within sample. Finally, this approach does not rely on any parametric assumptions on the return distributions of the two securities.

I then study three important applications of my approach. In the first application, I examine the cross-sectional equity premium associated with the sensitivity of stock returns to the market disaster risk. To capture this sensitivity, I construct a “systematic disaster risk” measure based on the conditional return distribution of a stock given the return of

the market proxied by the S&P500 index. For both the market and the individual stocks, I define a “normal” state and a “disaster” state. Then, the systematic disaster risk of each stock is defined as the difference in the conditional disaster probabilities of the stock given that the market is in the disaster versus the normal states, respectively. This measure captures the extent to which an individual stock is more likely to be hit by a rare disaster when the market moves from the normal state to the disaster state.

Intuitively, if a stock is more sensitive to the market disaster risk, then it should be less desirable for investors to hold, especially during time periods when a market crash is considered likely. As such, investors should require higher expected returns for holding stocks with higher systematic disaster risk, and this effect should be more pronounced when the market disaster risk is high. To test this hypothesis, I apply the Fama and MacBeth (1973) methodology. I find that systematic disaster risk is not priced when the option implied market disaster risk is low. However, when I restrict attention to time periods during which a market crash is perceived likely, then I find strong evidence that stocks with higher systematic disaster risk earn significantly higher expected returns after controlling for well documented risk factors. In fact, increasing the systematic disaster risk by one standard deviation raises expected monthly stock returns by 63 basis points, which is equivalent to over 7% per year.

My second application turns to the banking sector. I ask the question of how to estimate the exposure of banks to systemic shocks and what bank characteristics are related to banks’ systemic exposure. To this end, I construct a “systemic exposure” measure based on the conditional return distribution of a bank given the banking sector return, where the sector portfolio is empirically proxied by the KBW Bank Index. For both the sector portfolio and the individual banks, I define a “normal” state and a “disaster” state. I then estimate the systemic exposure of each bank as the difference in the conditional disaster probabilities of the bank given that the sector is in the disaster versus the normal states, respectively. Intuitively, the systemic exposure measure captures by how much an individual bank is more likely to experience a disaster when the whole banking sector falls from the normal state to the disaster state.

My estimates show that the systemic exposure measure is typically positive, indicating that banks are generally more likely to experience a disaster when the banking sector as

a whole is in the disaster state relative to when the sector performs normally. I also find that the systemic exposure of a bank increases with its equity beta and the total return volatility. It is also increasing in the non-interest to interest income ratio, reflecting that a bank's exposure to systemic shocks is largely driven by its non-traditional businesses. In addition, there is some evidence that systemic exposure decreases with total market capitalization.

Finally, in the third application I explore an extension of the recent Ross (2014) recovery theorem, which is aimed to recover the physical return distribution of the market portfolio from the corresponding risk-neutral distribution. While the recovery theorem deals with the market portfolio only, I seek to extend it to recover the physical return distribution of an individual asset. I show that this can be achieved through the risk-neutral joint return distribution of the asset with the market portfolio, which is given by the product of the risk-neutral marginal return distribution of the market and the conditional return distribution of the asset given the market return. Since my approach generates an estimate for this conditional return distribution, it lends itself naturally to the extension of the Ross recovery theorem to individual assets.

The rest of the paper proceeds as follows. Section 2 reviews the literature. Section 3 introduces the setup, linking the return distribution of an individual asset to that of an aggregate index. Section 4 discusses the estimation methodology and how it can be implemented using option prices. Section 5 provides some discussions and extensions. Section 6 applies the framework to examine the cross-sectional systematic disaster risk premium. Section 7 studies banks' exposure to systemic shocks. Section 8 shows how my approach can be used to extend the Ross recovery theorem to individual assets. Section 9 concludes. Proofs of propositions are shown in Appendix A, and other technical discussions are delegated to Appendix B.

2 Literature Review

The paper contributes to several strands of the literature. First, it adds to the study of the joint return behavior of different securities. Roll (1988), Jorion (2000), and Longin and Solnik (2001), among others, show that the return correlation between securities is

not symmetric under all market conditions, but instead increases during market crashes. Ang and Chen (2002), Hong, Tu and Zhou (2007), and Jiang, Wu and Zhou (2014) develop methods to test this asymmetric dependence between security returns. Skinzi and Refenes (2004) and Driessen, Maenhout, and Vilkov (2013) propose methods of inferring equity return correlations from option prices. In this paper, I provide a novel approach of estimating the entire conditional return distribution of an asset given the return of an aggregate index, thus accounting for potential asymmetries in their joint behavior. In addition, this approach captures all moment properties of their joint distribution, which is beyond the return correlation alone.

The paper also adds to the extensive literature on estimating the risk-neutral distributions of security returns using option pricing data. Ross (1976) and Breeden and Litzenberger (1978) first show that one can extract the risk-neutral probability distributions of security returns from option prices. Since option prices are available at discrete strike prices and maturities only, some smoothing techniques are needed to estimate the full risk-neutral distribution. (See Melick and Thomas (1997), Posner and Milevsky (1998), and Rubinstein (1998) for parametric methods and Shimko (1993), Jackwerth and Rubinstein (1996), Malz (1997), and Ait-Sahalia and Lo (1998) for non-parametric methods.) Jackwerth (1999) provides a comprehensive review on various methods used to extract the risk-neutral return distributions from option prices. Figlewski (2010) provides an empirical demonstration based on the U.S. market portfolio. Overall, the literature has restricted attention to the return distributions of single securities. My paper extends this literature by introducing an approach of estimating the conditional return distribution of an asset given the return of an aggregate index using option prices. Under the assumption that the variation in the index return is the only priced risk, the risk-neutral conditional distribution estimated from option prices coincides with the physical conditional distribution.

This paper is also related to research on equity premium associated with disaster risk. A large body of theoretical research shows that investors are averse to rare disasters (e.g., Barro (2006, 2009), Gabaix (2008, 2012), Gourio (2012), Chen, Joslin, and Tran (2012), and Wachter (2013)). Consistent with this, a number of empirical papers have documented a positive relation between disaster risk and expected market returns (e.g., Bali, Demirtas, and Levy (2009), Bollerslev and Todorov (2011), and Jiang and Kelly (2013)). Cross-

sectionally, Siriwardane (2013) studies the relation between expected asset returns and the option-implied disaster risk of assets, and finds a positive premium. Van Oordt and Zhou (2012), Jiang and Kelly (2013), Ruenzi and Weigert (2013), on the other hand, focus on the systematic portion of disaster risk by looking at the disaster risk of an asset in relation to that of the market. Given the challenge of estimating the joint disaster risk due to the rare occurrence of disastrous events, all three papers use historical equity returns and resort to either the power law distribution or parametric copulas to model the tails, which unfortunately does not necessarily represent the true probability distributions. My paper contributes to this literature by suggesting a measure of systematic disaster risk based on the conditional return distribution of an asset given the market return. Since the measure is estimated using option prices, it is forward-looking and naturally captures rare disasters. This measure reflects investors' perceived sensitivity of asset returns to the market disaster risk, which can be conveniently used to test the cross-sectional disaster risk premium.

This paper also adds to the literature on bank systemic risk. By definition, systemic risk focuses on risk associated with the collapse of the entire banking system. Hence, the main challenge of estimating systemic risk comes from the rare occurrence of disastrous events. Different methods have been proposed to tackle this problem. For example, Huang, Zhou, and Zhu (2009) measure systemic risk by the price of insurance against financial distress, in which the default correlation between banks is proxied by the equity return correlation. Acharya, Pedersen, Philippon and Richardson (2010) propose the systemic expected short-fall (*SES*) measure, which estimates the propensity of a bank to be undercapitalized when the system as a whole is undercapitalized. In particular, their method relies on the power law distribution to model the tails. Adrian and Brunnermeier (2011) propose the $\Delta CoVaR$ measure as the difference between the value-at-risk of the banking system conditional on an individual bank being in distress and the value-at-risk of the banking system conditional on the bank being solvent. Empirical estimation of $\Delta CoVaR$ uses quantile regression to capture the tail distributions. The contribution of my paper is that it provides a measure of banks' exposure to systemic shocks based on the conditional return distribution of a bank given the banking sector return. Since this measure is estimated using option prices, it is forward-looking and naturally captures investors' perceived exposure of a bank to sector-wide disastrous shocks even if such shocks do not occur within sample.

Finally, the paper also contributes to the literature on Ross (2014) recovery theorem, which recovers the physical probability distribution of the market return from the associated risk-neutral distribution. Subsequent research has been done to further explore this problem. Carr and Yu (2012) provide alternative assumptions that allow for recovery for diffusions on a bounded state space. Huang and Shaliastovich (2013) develop a recursive-utility framework to separately identify physical probabilities and risk adjustments. Martin and Ross (2013) show that recovery can indeed be effected by observing the behavior of the long end of the yield curve. Walden (2013) extends the Ross recovery result to unbounded diffusion processes. See also Dubynskiy and Goldstein (2013) and Borovička, Hansen and Scheinkman (2014) for criticism of the Ross recovery theorem. This literature primarily focuses on the market portfolio. My approach contributes to this literature by extending the recovery results to any individual asset through its risk-neutral joint return distribution with the market.

3 Setup

In this section, I introduce a simple setup that links the return distribution of an individual asset with that of an aggregate index through the total probability formula. The return of the aggregate index can be viewed as a factor that determines the state of the economy. Examples of the index include the market portfolio or a sector portfolio, etc. The next section will discuss how this setup allows me to estimate the conditional return distribution of the asset given the index return by a regression approach.

There are T discrete time points $t \in \{1, 2, \dots, T\}$. At any time t , I consider security returns over one period ahead, that is, from t to $t + 1$.

Consider an aggregate index I , whose return over any one period can take N values

$$(r^I(1), r^I(2), \dots, r^I(N)).$$

Denote by $\tilde{r}_{t,t+1}^I$ the random return of the index over the period from t to $t + 1$. Evaluated at time t , the probability distribution of $\tilde{r}_{t,t+1}^I$ is given by the vector

$$\mathbf{p}_{t,t+1}^I = (p_{t,t+1}^I(1), p_{t,t+1}^I(2), \dots, p_{t,t+1}^I(N)),$$

where $p_{t,t+1}^I(n)$ represents the probability of $\tilde{r}_{t,t+1}^I = r^I(n)$ for any $n \in \{1, 2, \dots, N\}$.

Consider an asset, whose return over any one period can take K distinct values

$$(r(1), r(2), \dots, r(K)).$$

Denote by $\tilde{r}_{t,t+1}$ the random return of the asset over the period from t to $t+1$. Evaluated at t , the probability distribution of $\tilde{r}_{t,t+1}$ is given by the vector

$$\mathbf{p}_{t,t+1} = (p_{t,t+1}(1), p_{t,t+1}(2), \dots, p_{t,t+1}(K)),$$

where $p_{t,t+1}(k)$ represents the probability of $\tilde{r}_{t,t+1} = r(k)$ for any $k \in \{1, 2, \dots, K\}$.

Assumption 1 (*Identifying*) *The conditional probability distribution of the asset return given any value of the contemporaneous index return does not vary over time.*

This assumption may be understood in relation to the one we make when empirically estimating the CAPM beta that the conditional mean return of an asset given any value of the market return is fixed over time. My assumption is stronger in the sense that it requires not only the conditional mean but indeed the entire conditional distribution to be time-invariant. It implies that the time variation in the return distribution of the asset ($\mathbf{p}_{t,t+1}$) is solely driven by the time variation in the return distribution of the index ($\mathbf{p}_{t,t+1}^I$).

Denote the time-invariant conditional distribution of the asset return given the index return by the matrix

$$\boldsymbol{\theta} = \begin{pmatrix} \theta(1|1) & \cdots & \theta(K|1) \\ \vdots & \ddots & \vdots \\ \theta(1|N) & \cdots & \theta(K|N) \end{pmatrix},$$

where $\theta(k|n)$ stands for the conditional probability of $\tilde{r}_{t,t+1} = r(k)$ given $\tilde{r}_{t,t+1}^I = r^I(n)$ evaluated at the beginning of the period for any $n \in \{1, 2, \dots, N\}$ and $k \in \{1, 2, \dots, K\}$. By Assumption 1, $\boldsymbol{\theta}$ does not depend on time. According to the properties of conditional probabilities, it must be that given any value of the index return, the conditional probabilities of the asset return sum up to one, i.e.,

$$\sum_{k=1}^K \theta(k|n) = 1, \quad \forall n. \tag{1}$$

At any t , the marginal return distribution of the asset is related to that of the index by the total probability formula. Specifically, for any k ,

$$p_{t,t+1}(k) = \sum_{n=1}^N p_{t,t+1}^I(n) \theta(k|n).$$

That is, the marginal distribution of the asset return is equal to the weighted average of its conditional distribution given the index return, with the weights given by the marginal return distribution of the index. This can be conveniently rewritten in matrix form as

$$\mathbf{p}_{t,t+1} = \mathbf{p}_{t,t+1}^I \cdot \boldsymbol{\theta}. \quad (2)$$

The discussion so far is based on the physical probability measure. I now make an additional assumption to link the physical measure with the risk-neutral measure.

Assumption 2 *The variation in the index return is the only priced (systematic) risk.*

Formally, Assumption 2 is satisfied if there exists a stochastic discount factor, whose value related to future payoffs depends on the future value of the index return only. This assumption can be understood in relation to the CAPM framework, which implies that the stochastic discount factor is a linear function of the market return. Assumption 2 is weaker in the sense that it requires the stochastic discount factor to be a function of the index return only, but it does not impose any restriction on the functional form of this relation. This assumption implies that conditional on a particular value of the index return, any variation left in the asset return is purely idiosyncratic and hence is risk-neutrally priced. It is then straightforward to show that the conditional distribution of the asset return given any value of the index return is the same under both the physical and the risk-neutral probability measures. This is formally stated in the following proposition.

Proposition 1 *At any time t , the risk-neutral conditional probability of $\tilde{r}_{t,t+1} = r(k)$ given $\tilde{r}_{t,t+1}^I = r^I(n)$ is equal to $\theta(k|n)$ for all n and k .*

A conclusion of Proposition 1 is that the total probability formula (2) holds just as well under the risk-neutral probability measure with respect to the same conditional probability

matrix $\boldsymbol{\theta}$. At time t , denote the risk-neutral distributions of $\tilde{r}_{t,t+1}$ and $\tilde{r}_{t,t+1}^I$ by

$$\begin{aligned}\mathbf{q}_{t,t+1}^I &= (q_{t,t+1}^I(1), q_{t,t+1}^I(2), \dots, q_{t,t+1}^I(N)), \\ \mathbf{q}_{t,t+1} &= (q_{t,t+1}(1), q_{t,t+1}(2), \dots, q_{t,t+1}(K)).\end{aligned}$$

Formally,

Corollary 1 *At any time t ,*

$$\mathbf{q}_{t,t+1} = \mathbf{q}_{t,t+1}^I \cdot \boldsymbol{\theta}. \quad (3)$$

4 Estimation Methodology

The setup introduced in Section 3 can be used to estimate the conditional distribution matrix $\boldsymbol{\theta}$ from the marginal return distributions of the two securities. Proposition 1 and Corollary 1 suggest that for this purpose one can work under either the physical measure or the risk-neutral measure, and the conditional probabilities obtained under both measures would be identical. In practice, there can be different ways of estimating the marginal return distributions in either probability measure. In this paper, I choose to do so under the risk-neutral measure using option prices based on the work of Ross (1976) and Breeden and Litzenberger (1978). My approach would work in the same manner if the marginal return distributions are obtained using other methods.

I assume that both index I and the individual asset of interest are traded in the option market. Examples of aggregate indices with traded options include the S&P500 index and the KBW Bank Index, etc. The procedure of estimating $\boldsymbol{\theta}$ includes two steps. In the first, I extract the risk-neutral marginal return distributions of the index and the asset, $\mathbf{q}_{t,t+1}^I$ and $\mathbf{q}_{t,t+1}$, from option prices. Then in the second step, I perform a constrained regression of $\mathbf{q}_{t,t+1}$ on $\mathbf{q}_{t,t+1}^I$ over time to estimate the conditional distribution matrix $\boldsymbol{\theta}$. Below I discuss each of the two steps separately.

4.1 Extracting Risk-Neutral Marginal Distributions from Option Prices

I first discuss how the risk-neutral return distributions $\mathbf{q}_{t,t+\tau}^I$ and $\mathbf{q}_{t,t+\tau}$ can be extracted from option prices. The estimation procedures for $\mathbf{q}_{t,t+1}^I$ and $\mathbf{q}_{t,t+1}$ are parallel, and hence in this section I focus on $\mathbf{q}_{t,t+1}$ for brevity.

Ross (1976) and Breeden and Litzenberger (1978) show that given a continuous range of strike prices covering all possible values of the underlying asset at maturity, the entire risk-neutral probability distribution of the asset's future value can be estimated from European option prices. Suppose that t is the current time point and consider an European put option that matures at time $t + 1$. Let S_t represent the current price of the underlying asset, and let \tilde{S}_{t+1} be the random price of the asset in one period. Denote the strike price of the option by X and the risk-free rate by r_f . The price of the put option can then be expressed as a function of the strike price:

$$\begin{aligned} Put(X) &= e^{-r_f} \int_0^\infty (X - S_{t+1})^+ dF(S_{t+1}) \\ &= e^{-r_f} \int_0^X (X - S_{t+1}) dF(S_{t+1}), \end{aligned} \tag{4}$$

where $F(\cdot)$ is the risk-neutral cumulative distribution function (CDF) of \tilde{S}_{t+1} evaluated at time t . Differentiating (4) with respect to X obtains

$$\frac{\partial Put(X)}{\partial X} = e^{-r_f} F(X).$$

Solving for $F(X)$ leads to

$$F(X) = e^{r_f} \frac{\partial Put(X)}{\partial X}. \tag{5}$$

Evaluating $F(X)$ at all possible values of X thus yields the entire risk-neutral distribution of the asset price at maturity.¹

Since I am interested in the risk-neutral distribution of the asset return from time t to $t + 1$, I need to relate the return to the price at maturity. Assume that the dividend yield paid by the asset from t to $t + 1$ is equal to d . Then, the return of the asset $\tilde{r}_{t,t+1}$ is related to the future asset price \tilde{S}_{t+1} according to the following approximation

$$\tilde{S}_{t+1} = S_t (1 + \tilde{r}_{t,t+1} - d). \tag{6}$$

Therefore, the risk-neutral CDF of the asset return is given by $\forall r$,

$$G(r) = F(S_t (1 + r - d)).$$

¹One can alternatively estimate $F(\cdot)$ based on prices of European call options. By the call-put parity, the results using call and put options are identical. For simplicity, I choose to work with put options.

Section 3 assumed that the asset return took a finite number (K) of values. This, however, is a simplification of the real world in which asset returns have continuous ranges. To be consistent with the setup, I discretize the continuous asset return by dividing its range into K mutually disjoint intervals with $K - 1$ thresholds. At each time t , the risk-neutral probabilities that the one-period asset return lies in each of these K intervals are taken as elements of the vector $\mathbf{q}_{t,t+1}$. In particular, let $rr(1), rr(2), \dots, rr(K-1)$ denote the $K - 1$ thresholds separating the K intervals of the asset return. Then, the vector $\mathbf{q}_{t,t+1}$ can be estimated as

$$\begin{aligned} q_{t,t+1}(1) &= G(rr(1)), \\ q_{t,t+1}(k) &= G(rr(k)) - G(rr(k-1)), \quad \forall k = 2, 3, \dots, K-1, \\ q_{t,t+1}(K) &= 1 - G(rr(K-1)). \end{aligned}$$

Two additional technical issues need to be dealt with for the empirical estimation of $\mathbf{q}_{t,t+1}$. First, the estimation of the risk-neutral CDF (5) relies on differentiating the option price with respect to the strike price. Since it is generally very difficult to obtain a close-form expression for this derivative, I estimate it by linear approximation. A second empirical challenge has to do with obtaining European option prices. Nearly all individual stock options are American options. While indices are generally represented by European options, the market option prices are only available at discrete values of the strike price and time to maturity. To obtain the European option price for any security at any arbitrary point, I adopt a simple and commonly used approach of first fitting the implied volatility surface by kernel smoothing and then deriving the Black-Merton-Scholes (BMS) option price (Black and Scholes (1973) and Merton (1973)) using the fitted volatility. I delegate detailed discussions of these issues to Appendix B.

4.2 Estimating Conditional Distributions by Constrained Regression

This section discusses how to estimate the conditional distribution matrix θ based on the marginal return distributions $\mathbf{q}_{t,t+\tau}^I$ and $\mathbf{q}_{t,t+\tau}$. Since $\mathbf{q}_{t,t+\tau}^I$ and $\mathbf{q}_{t,t+\tau}$ are extracted from option prices, they are often subject to measurement errors. In particular, since options written on individual assets are more thinly traded than index options, one would expect $\mathbf{q}_{t,t+\tau}$ to be much noisier than $\mathbf{q}_{t,t+\tau}^I$. To reflect the different degrees of noisiness in $\mathbf{q}_{t,t+\tau}^I$

and $\mathbf{q}_{t,t+\tau}$, I assume that $\mathbf{q}_{t,t+\tau}^I$ can be accurately estimated and that $\mathbf{q}_{t,t+\tau}$ contains noises, which are captured by an error term

$$\boldsymbol{\epsilon}_{t,t+1} = (\epsilon_{t,t+1}(1), \epsilon_{t,t+1}(2), \dots, \epsilon_{t,t+1}(K)).$$

Now the risk-neutral total probability formula (3) becomes

$$\mathbf{q}_{t,t+1} = \mathbf{q}_{t,t+1}^I \cdot \boldsymbol{\theta} + \boldsymbol{\epsilon}_{t,t+1}. \quad (7)$$

I will estimate $\boldsymbol{\theta}$ from $\mathbf{q}_{t,t+\tau}^I$ and $\mathbf{q}_{t,t+\tau}$ using a constrained linear regression based on (7). Before discussing the detailed estimation procedures, I need to make some additional assumptions, which are sufficient to maintain the consistency of my estimates. In particular, I assume the following.

Assumption 3 *The pair of vectors $\{\mathbf{q}_{t,t+1}^I, \mathbf{q}_{t,t+1}\}$ are jointly stationary and weakly dependent over time.²*

Assumption 4 *At any time t , $E[\epsilon_{t,t+1}(k) | \mathbf{q}_{t,t+1}^I] = 0$ for all $k \in \{1, 2, \dots, K\}$, where the expectation is taken under the physical probability measure.*

Assumption 5 *The $T \times N$ matrix*

$$\mathbf{Q}^I = \begin{pmatrix} \mathbf{q}_{1,2}^I \\ \mathbf{q}_{2,3}^I \\ \vdots \\ \mathbf{q}_{T,T+1}^I \end{pmatrix}$$

is of rank N .

To see how $\boldsymbol{\theta}$ can be estimated by linear regression, it is useful to rewrite (7) as

$$q_{t,t+1}(k) = \sum_{n=1}^N q_{t,t+1}^I(n) \theta(k|n) + \epsilon_{t,t+1}(k), \quad \forall k.$$

This indicates that one can estimate $(\theta(k|1), \theta(k|2), \dots, \theta(k|N))'$ (the k^{th} column of the conditional distribution matrix $\boldsymbol{\theta}$) by running an OLS regression of $q_{t,t+1}(k)$ on the vector

²The pair of vectors $\{\mathbf{q}_{t,t+1}^I, \mathbf{q}_{t,t+1}\}$ are weakly dependent over time if for any t , $\{\mathbf{q}_{t,t+1}^I, \mathbf{q}_{t,t+1}\}$ and $\{\mathbf{q}_{t+\Delta t, t+\Delta t+1}^I, \mathbf{q}_{t+\Delta t, t+\Delta t+1}\}$ become approximately independent as $\Delta t \rightarrow \infty$.

$\mathbf{p}_{t,t+1}^I$ over time. Assumptions 3–5 guarantee that the resulting OLS estimates are consistent, i.e., they converge to the true parameter values when the sample size approaches infinity.

Then, to estimate the entire $\boldsymbol{\theta}$ matrix, a natural idea would be to run a total of K regressions corresponding to each of the K values of the asset return. However, since $\boldsymbol{\theta}$ represents the conditional probabilities, two implicit constraints must be satisfied. First is that the conditional probabilities given any value of the index return must sum up to one (as required by (1)), and the second constraint says that all elements of $\boldsymbol{\theta}$ must lie between 0 and 1, i.e., $0 \leq \theta(k|n) \leq 1$ for all k and n . Unfortunately, running K OLS regressions independently does not guarantee that these constraints are satisfied.

A solution to this issue is to conduct the K regressions jointly subject to the above two constraints. Formally, define

$$\mathbf{Q} = \begin{pmatrix} \mathbf{q}_{1,2} \\ \mathbf{q}_{2,3} \\ \vdots \\ \mathbf{q}_{T,T+1} \end{pmatrix},$$

$$\boldsymbol{\epsilon} = \begin{pmatrix} \epsilon_{1,2}(1) & \cdots & \epsilon_{1,2}(K) \\ \vdots & \ddots & \vdots \\ \epsilon_{T,T+1}(1) & \cdots & \epsilon_{T,T+1}(K) \end{pmatrix},$$

and let $\mathbf{1}_{a \times b}$ and $\mathbf{0}_{a \times b}$ denote the $a \times b$ matrices of ones and zeros for any positive integers a and b , respectively. Then, the problem can be represented by the following constrained linear regression

$$\mathbf{Q} = \mathbf{Q}^I \cdot \boldsymbol{\theta} + \boldsymbol{\epsilon}, \tag{8}$$

s.t.

$$\boldsymbol{\theta} \cdot \mathbf{1}_{K \times 1} = \mathbf{1}_{N \times 1},$$

$$\mathbf{0}_{N \times K} \leq \boldsymbol{\theta} \leq \mathbf{1}_{N \times K}.$$

I denote the resulting estimates from problem (8) by $\hat{\boldsymbol{\theta}}_T$, where the subscript T reflects dependence of the estimates on the sample size.

A priori, it is not clear whether imposing the constraints would affect the consistency of my estimation. The following proposition establishes that consistency is indeed preserved in the presence of the constraints.

Proposition 2 *Under Assumptions 3-5, $\hat{\theta}_T$ is a consistent estimator of θ , i.e., $\lim_{T \rightarrow \infty} \hat{\theta}_T = \theta$.*

5 Discussions and Extensions

This section provides some discussions and extensions of the estimation framework introduced above.

5.1 Elaboration on Key Assumptions

Assumptions 1 and 2 are key to my approach in that they point out two important roles of the index return.

Assumption 1 states that the conditional return distribution of the asset given the index return is time invariant. The intuition is that while the return distribution of the asset can change over time, its variation is solely driven by the time variation in the distribution of the index return. In particular, once the index return is fixed, the conditional distribution of the asset return is also fixed, regardless of the time point under consideration. This is the identifying assumption of my approach in that it allows me to make use of the time series information on the marginal return distributions of the two securities to determine the time-invariant conditional probabilities. Without this assumption, my estimation model is not identified.

Another important role of the index return is reflected in Assumption 2, which states that the variation in the index return is the only priced (systematic) risk. As such, fixing a certain value of the index return, any variation left of the asset return is purely idiosyncratic and is thus risk-neutrally priced. This assumption implies that the conditional return distribution of the asset given the index return is the same under the physical and the risk-neutral probability measures. Since investors are averse to risk, security return distributions generally differ under the physical versus the risk neutral measures to reflect the adjustment for risk aversion. In fact, there is a recent literature on Ross (2014) recovery theorem that aims to recover the physical return distributions from the associated risk-neutral distributions. The benefit of Assumption 2 is that it aligns the analyses under the two probability measures once I condition on a particular value of the index return. This

allows me to perform empirical estimation under the risk-neutral measure using option prices, and the resulting conditional probabilities would be exactly the same as if I work under the physical measure. However, the failure of this assumption does not necessarily invalidate my approach. Even when this assumption is violated (e.g., when the Fama-French three-factor pricing model holds), my approach can still be applied under either the physical or the risk-neutral measure, but the conditional distribution of the asset return would no longer be the same under the two probability measures.

5.2 Multi-Factor Framework

The baseline model discussed earlier is a one-factor framework, in which the index return is the only factor that determines the asset return distributions. In practice, asset return distributions can be affected by more than one factor. Macroeconomic variables such as the consumption growth rate, inflation rate, or VIX may serve as additional factors. In this case, I need to extend my model into a multi-factor framework.

Suppose that there exist M factors. Each factor $m \in \{1, 2, \dots, M\}$ takes N^m values. Evaluated at time t , the joint distribution of the M factors at time $t + 1$ is given by the joint distribution function $p_{t,t+1}^I(n^1, n^2, \dots, n^M)$, where $n^m \in \{1, 2, \dots, N^m\}$ for every m . Similar to Assumption 1, I assume that given any set of joint values of these M factors, the conditional return distribution of an asset does not vary over time, which is denoted by the conditional distribution function $\theta(\cdot | n^1, n^2, \dots, n^M)$. Then, the total probability formula links the marginal return distribution of the asset to the joint distribution of the M factors through $\theta(\cdot | n^1, n^2, \dots, n^M)$, i.e.,

$$p_{t,t+1}(k) = \sum_{n^1=1}^{N^1} \sum_{n^2=1}^{N^2} \dots \sum_{n^M=1}^{N^M} p_{t,t+1}^I(n^1, n^2, \dots, n^M) \theta(k | n^1, n^2, \dots, n^M).$$

If I further assume that variations in these M factors constitute the only priced (systematic) risk (the multi-factor version of Assumption 2), then $\theta(\cdot | n^1, n^2, \dots, n^M)$ is the same under the physical and the risk-neutral measures. This allows me to write down the risk-neutral total probability formula as

$$q_{t,t+1}(k) = \sum_{n^1=1}^{N^1} \sum_{n^2=1}^{N^2} \dots \sum_{n^M=1}^{N^M} q_{t,t+1}^I(n^1, n^2, \dots, n^M) \theta(k | n^1, n^2, \dots, n^M),$$

where $q_{t,t+1}^I(n^1, n^2, \dots, n^M)$ represents the risk-neutral joint distribution of the M factors at $t + 1$ evaluated at time t .

If I have the marginal return distribution of the asset and the joint distribution of the factors under either the physical or the risk neutral measure, I can estimate the conditional distribution function $\theta(k|n^1, n^2, \dots, n^M)$ by regressing the former on the latter over time, as in the baseline framework. Unfortunately, the risk-neutral joint distribution of the factors can no longer be directly extracted from option prices, because options written on the joint values of multiple factors are generally not available. As a result, one need to resort to other approaches to obtain the joint distribution of the factors. Once this joint distribution is obtained, I can estimate $\theta(k|n^1, n^2, \dots, n^M)$ in exactly the same manner as in the baseline case.

5.3 Continuous Security Returns

Up till now, I have assumed that the returns of both the individual asset and the aggregate index take a finite number of discrete values. This section considers the case of continuous security returns. Given the analogy between the physical and risk-neutral analyses (as in the baseline case), in this section I work directly with the risk-neutral measure.

At any time t , suppose that the one-period index return $\tilde{r}_{t,t+1}^I$ and the one-period asset return $\tilde{r}_{t,t+1}$ take continuous values from the interval $[-1, \infty)$. The marginal probability distributions of $\tilde{r}_{t,t+1}^I$ and $\tilde{r}_{t,t+1}$ are given by the density functions $q_{t,t+1}^I(\cdot)$ and $q_{t,t+1}(\cdot)$, respectively. By Assumption 1, given any value of $\tilde{r}_{t,t+1}^I$ the conditional distribution of $\tilde{r}_{t,t+1}$ does not change over time. I denote this time-invariant conditional distribution by the conditional density function $\theta(\cdot|\cdot)$, which integrates to 1 given any $\tilde{r}_{t,t+1}^I = r^I$, i.e.,

$$\int_{-1}^{\infty} \theta(r|r^I) dr = 1, \quad \forall r^I.$$

At any t , the marginal return distributions of the two securities are linked to each other by the total probability formula

$$q_{t,t+1}(r) = \int_{-1}^{\infty} q_{t,t+1}^I(r^I) \theta(r|r^I) dr^I.$$

Assume that $q_{t,t+1}^I(\cdot)$ can be accurately measured, whereas $q_{t,t+1}(\cdot)$ is subject to noises,

which are captured by the error term $\epsilon_{t,t+1}(\cdot)$. Then, the total probability formula becomes

$$q_{t,t+1}(r) = \int_{-1}^{\infty} q_{t,t+1}^I(r^I) \theta(r|r^I) dr^I + \epsilon_{t,t+1}(r).$$

Since $\theta(r|r^I)$ is infinite-dimensional, its empirical estimation is difficult without further information on the structure of $\theta(r|r^I)$. While there are different ways of reducing dimensionality, one of the simplest methods is to make parametric assumptions on the functional form of $\theta(r|r^I)$. Specifically, let

$$\theta(r|r^I) = g(r, r^I; \boldsymbol{\lambda}),$$

where $g(r, r^I; \boldsymbol{\lambda})$ is the assumed functional form of $\theta(r|r^I)$ and $\boldsymbol{\lambda}$ represents the parameters of choice. Then, one can estimate $\theta(r|r^I)$ by choosing the values of $\boldsymbol{\lambda}$ to solve the following least square problem:

$$\begin{aligned} & \min_{\boldsymbol{\lambda}} \left(\sum_{t=1}^T \int_{-1}^{\infty} \left[q_{t,t+1}(r) - \int_{-1}^{\infty} q_{t,t+1}^I(r^I) g(r, r^I; \boldsymbol{\lambda}) dr^I \right]^2 dr \right), \\ & \text{s.t.} \\ & \int_{-1}^{\infty} g(r, r^I; \boldsymbol{\lambda}) dr = 1, \quad \forall r^I, \\ & g(r, r^I; \boldsymbol{\lambda}) \geq 0, \quad \forall r^I, r. \end{aligned}$$

5.4 Alternative Econometric Models

In Section 4.2, I used a constrained linear regression to estimate the conditional distribution matrix $\boldsymbol{\theta}$ from the marginal return distributions of the two securities. In fact, there are some alternative econometric models (rather than the constrained linear regression) that may also seem appealing for my purpose. I discuss the potentials and limitations of some alternatives in this section.

Probit and Logit Models

The Probit and Logit models both can be used to predict the probability distribution of an outcome variable based on the values of the independent variables. The two models differ in the assumed distribution of the error term. The benefit of these models is that they automatically guarantee that the estimated probabilities of the outcome variable lie

between zero and one. It may seem that the Probit and the Logit models are well suited for my purpose. However, a key difference is that in these two models, the dependent variable is a discrete variable representing the outcome of an event. In contrast, in my case the dependent variable itself is the probability distribution of the asset return. Therefore, the Probit and the Logit models do not apply here. In addition, to obtain reasonable results I impose two constraints on my estimates, requiring that each conditional probability lie between zero and one and that they sum up to one given any particular value of the index return. It is not trivial to incorporate these constraints into the Probit and Logit models. In fact, it is not hard to see that once these constraints are met, the predicted marginal probabilities of the asset return based on the current linear model are guaranteed to lie between zero and one with no need for additional restrictions.

Maximum Likelihood Estimation

Another alternative econometric approach worth mentioning is the Maximum Likelihood Estimation (MLE). Suppose that the joint distribution of the error terms $\boldsymbol{\epsilon}_{t,t+1} = (\epsilon_{t,t+1}(1), \epsilon_{t,t+1}(2), \dots, \epsilon_{t,t+1}(K))$ is given by the joint density function $\Lambda(\cdot)$. If $\boldsymbol{\epsilon}_{t,t+1}$ is independent and identically distributed over time, then the conditional distribution matrix $\boldsymbol{\theta}$ can be estimated by the following constrained MLE:

$$\begin{aligned} & \max \prod_{t=1}^T \Lambda(\mathbf{q}_{t,t+1} - \mathbf{q}_{t,t+1}^I \cdot \boldsymbol{\theta}), \\ & s.t. \\ & \boldsymbol{\theta} \cdot \mathbf{1}_{K \times 1} = \mathbf{1}_{N \times 1}, \\ & \mathbf{0}_{N \times K} \leq \boldsymbol{\theta} \leq \mathbf{1}_{N \times K}. \end{aligned}$$

The key here is the joint density function $\Lambda(\cdot)$. It is not clear what the best assumption would be for the joint distribution of $\boldsymbol{\epsilon}_{t,t+1}$. The normal distribution, for instance, may not be a good choice. This is because both the true and the estimated marginal return distributions of the asset ($\mathbf{q}_{t,t+1}$) have bounded values, and hence the associated measurement errors $\boldsymbol{\epsilon}_{t,t+1}$ should also be bounded, which is clearly not the case for normally distribution variables. In addition, it is also likely that the different elements of $\boldsymbol{\epsilon}_{t,t+1}$ are correlated with each other, rendering the assumption on $\Lambda(\cdot)$ even more complicated.

While the constrained linear regression model adopted in this paper seems simple, I will provide evidence for the out-of-sample validity of my estimation in the applications to be discussed in the following sections.

6 Application I: Systematic Disaster Risk Premium

Starting from the seminal work of the Capital Asset Pricing Model (CAPM), independently developed by Sharpe (1964), Lintner (1965a,b), and Mossin (1966), researchers have found that the cross-sectional risk-return relation is driven by the comovement of individual asset returns with the market return, which is usually termed “systematic risk.” Rubinstein (1973) and Kraus and Litzenberger (1976) extend the CAPM framework, which focuses on the second moment of security returns, to account for higher moments. More recently, Kadan, Liu, and Liu (2014) propose a general framework of evaluating systematic risk for a broad class of risk measures, potentially accounting for various risk attributes such as high distribution moments, downside risk and rare disasters.

On the other hand, a large body of theoretical research shows that investors are averse to rare disasters (e.g., Barro (2006, 2009), Gabaix (2008, 2012), Gourio (2012), Chen, Joslin, and Tran (2012), and Wachter (2013)). Consistent with this, a number of empirical papers have documented a positive relation between disaster risk and expected market returns (e.g., Bali, Demirtas, and Levy (2009), Bollerslev and Todorov (2011), and Jiang and Kelly (2013)).

Given that investors exhibit aversion to rare disasters and that they are concerned with the comovement of asset returns with the market, it is then natural to conjecture that investors require higher compensation for holding assets that are more sensitive to the market disaster risk. This idea is empirically tested in the literature by Van Oordt and Zhou (2012), Jiang and Kelly (2013), and Ruenzi and Weigert (2013). All three papers show that stocks with higher sensitivity to the market disaster risk earn higher expected returns, at least during some time period. Given the challenge of estimating disaster risk due to the rare occurrence of disastrous events, all of these papers use historical returns and model the tails by either the power law distribution or parametric copulas.

In this section, I propose a “systematic disaster risk” measure, which captures the

sensitivity of asset returns to the market disaster risk. The measure is constructed based on the conditional return distribution of an asset given the market return, which can be readily estimated using my approach by taking the market portfolio as the related index. My measure is estimated using option prices rather than historical security returns. Thus, it is forward-looking and naturally accounts for investors' beliefs on the likelihood of rare events even if such events do not occur within sample. I empirically examine whether this measure predicts the cross-section of expected stock returns.

6.1 Measure of Systematic Disaster Risk

I construct the systematic disaster risk measure of a stock based on its conditional return distribution given the market return. For both the stocks and the market portfolio, I focus on three-month returns, whose associated risk-neutral distributions can be extracted from prices of options maturing in three months. The reason for using three-month returns is because options with a maturity around three months are considered to have the most accurate prices.

For both the market portfolio and the individual stocks, I define two states ($N = K = 2$), the normal state (H) and the disaster state (L). I choose the disaster thresholds for the three-month returns of the market and individual stocks to be $-1/3$ and $-1/2$, respectively. In words, the market (a stock) experiences a disaster if it loses more than one third (one half) of its original value within three months. The disaster threshold for individual stocks is chosen to be lower than that for the market portfolio, because individual stock returns are more volatile than the market return. I will later show that the choice of the disaster thresholds is consistent with our notion of rare disasters.

For each stock, I define the systematic disaster risk measure as

$$SysDis = \theta(L|L) - \theta(L|H), \quad (9)$$

where $\theta(L|L)$ and $\theta(L|H)$ stand for the disaster probabilities of the stock conditional on the market being in the disaster and normal states, respectively. Intuitively, this measure captures the extent to which an individual stock is more likely to experience a disaster when the whole market moves from the normal state to the disaster state. In other words, it is a measure of the sensitivity of individual stock returns to the market disaster risk.

Since $\theta(L|L)$ and $\theta(L|H)$ are conditional probabilities and take values from 0 to 1, it is apparent that the systematic disaster risk measure $SysDis$ ranges from -1 to 1. In particular, a positive value means that a stock is more likely to have a disaster when the market as a whole is in the disaster state relative to when the market performs normally. The higher the value of $SysDis$, the more sensitive the stock is to the market disaster risk. On the other hand, a negative $SysDis$ implies that the stock is less likely to experience a disaster when the market falls into the disaster state, and therefore can be viewed as a hedge against the market disaster risk.

6.2 Data

My sample period is from January 1996 to December 2012. I start from year 1996 because option pricing information is not available before then. I use the S&P500 index as a proxy for the market portfolio. This index has actively traded call and put options that cover a wide range of moneyness and time to maturity levels. I focus on all firms traded on NYSE, AMEX, and NASDAQ with option prices available. My primary source of data is the OptionMetrics database, which contains multiple data sets and provides me with the following information for both the S&P500 index and the individual stocks.

- The *option prices* data set contains daily records of the implied volatility and option vega for all available options with a variety of strike prices and expiration dates. I use this information to fit the implied volatility surface by kernel smoothing. (See Appendix B for details.)
- The *security prices* data set has daily security prices and returns.
- The *index dividend yield* data set provides me with annualized dividend yield for the S&P500 index.
- The *dividend distribution history* data set has the dates and amounts of dividend payments for individual stocks. On each date t , I estimate the dividend yield for each stock over the three-month period ahead as the ratio of the total dividend payments during the period to the stock price on date t .

- The *zero coupon yield curve* data set contains continuously compounded zero-coupon interest rates for various numbers of days to maturity. The zero-coupon rate is used as a proxy for the risk-free rate in the calculation of the BMS option prices. Since I focus on prices of options with a maturity of three months, I need the three-month interest rate. When the interest rate with this maturity level is not directly given in the data set, I estimate it by linear approximation using rates for adjacent numbers of days to maturity available.

In addition to the above information from OptionMetrics, I also collect the following data. First, I obtain monthly returns for all stocks during the sample period from the Center for Research in Security Prices (CRSP) database. These monthly returns are used as the dependent variable in the cross-sectional test. In addition, to control for the CAPM beta in my test, I obtain daily excess returns of the CRSP value-weighted market portfolio from the Kenneth French online data library, which are useful for the estimation of beta. Finally, to control for the “value” effect documented in Fama and French (1993), I include the book-to-market ratio as a risk factor in the test. For this sake, I obtain the book value of equity for all stocks from the Compustat database.

6.3 Empirical Strategy

My hypothesis states that stocks with higher systematic disaster risk earn higher expected returns. In addition, I conjecture that this relation should be more pronounced when the market disaster risk becomes higher, i.e., when a market crash is considered to be more likely to happen. The intuition is as follows. Suppose that the market performs well and that the probability of a market crash is very low. Then, even if a stock could be sensitive to the market disaster risk, investors may not be much concerned given that a market disaster is very unlikely to happen. On the other hand, when investors believe that the probability of a market crash is considerably high, they may have strong incentives to avoid investing in stocks with high systematic disaster risk. This is because it now becomes very likely that investments in these stocks could incur large losses upon the occurrence of a market crash.

To test these hypotheses, I follow the standard approach of Fama and MacBeth (1973) and Fama and French (1992). For each month within the sample period, I perform a cross-

sectional regression of monthly excess stock returns on the lagged values of various stock characteristics. I then test the risk premium associated with a certain characteristic by conducting a t -test on the corresponding coefficient estimates across all months. To see how the systematic disaster risk premium depends on the market condition, I define a market disaster risk variable Dis^M , estimated annually as the average option-implied risk-neutral disaster probability of the market portfolio across all dates of the year. Apparently, a high value of Dis^M indicates that investors' perceived likelihood of a market crash is high. I then repeat the test using subsamples constructed based on the lagged value of Dis^M , and see how the relation between systematic disaster risk and expected stock returns changes across different subsamples.

The main stock characteristic of interest is the systematic disaster risk $SysDis$ defined in Section 6.1. I estimate this variable for each stock on an annual basis in three steps. First, on each date, I compute the option-implied risk-neutral probabilities of disaster for both the S&P500 index and each individual stock. For stocks with a low level of option trading around the disaster threshold, the limited availability of option prices can potentially affect the accuracy of estimation. To enhance accuracy, I compute the risk-neutral disaster probability for a stock on a date only when (1) there are at least ten different options written on the stock with pricing information available on that date, (2) the lowest moneyness level (ratio of strike price to stock price) of available options for the stock on that date is no higher than 0.95, and (3) the shortest (longest) time to maturity of available options for the stock on that date is lower (higher) than three months. In the second step, I run the constrained linear regression (8) for each stock-year using all dates within the year to obtain the conditional disaster probabilities of the stock given different states of the market. To improve accuracy, I perform this step only for stocks with risk-neutral disaster probability available for no less than 200 days during the year. Finally, I compute $SysDis$ using the definition (9).

To disentangle the effect of the systematic disaster risk on the expected stock return from that of the “unconditional” propensity of a stock to experience a disastrous event, I control for each stock's unconditional disaster risk (Dis) in the cross-sectional regressions. The Dis variable is estimated annually for each stock as the average option-implied risk-neutral disaster probability of the stock across all dates during the entire year.

To further control for other risk factors, I include the following stock characteristic variables in the test:

- CAPM beta (*Beta*), estimated annually using daily equity returns. Following Amihud (2002), I estimate beta by size portfolios. At the end of each year, I rank all stocks into ten size portfolios based on NYSE breakpoints.³ For each size portfolio, I compute the portfolio return as the equal-weighted average return of all stocks in the portfolio, based on which I then estimate the portfolio beta. Finally, I assign the portfolio beta to all stocks in the portfolio.
- Coskewness (*Coskew*), estimated annually using daily equity returns. Following Harvey and Siddique (2000), I measure coskewness as the covariance of stock excess returns and the square of market excess returns.
- Firm size (*Size*), estimated monthly as the logarithm of the total market capitalization.
- Book-to-market ratio (*B2M*), estimated annually as the ratio of the end-of-year book value of equity to the end-of-year market capitalization.
- Lagged stock return, proxied by the stock return (*Ret*) of the previous month. This is to capture the momentum effect documented in Jegadeesh and Titman (1993).
- Mean-adjusted illiquidity (*IlliqMA*), estimated annually for each stock according to Amihud (2002) as the annual average of the ratio of absolute daily return to daily dollar volume of trading, adjusted by the cross-sectional mean.

Notice that the variables *SysDis*, *Dis*, *Beta*, *Coskew*, *B2M*, and *IlliqMA* are all estimated on an annual basis. Therefore, in the cross-sectional regressions, I always regress excess stock returns in a particular month on the values of these variables estimated from the previous year. On the other hand, since *Size* and *Ret* are available monthly, I use the one-month lagged values of these two variables in the cross-sectional regressions to capture the size and momentum effects.

³As pointed out in Fama and French (1992), the idea of using NYSE breakpoints is to avoid having most size portfolios dominated by small stocks traded on NASDAQ.

6.4 Results

To make sure that my estimation of *SysDis* generates reasonable results, I first check the out-of-sample validity by trying to predict future states of stocks using my conditional disaster probability estimation. In particular, for each year I take the first quarter and try to predict the state of a stock over the quarter based on the realized state of the market during the same quarter and the conditional disaster probabilities estimated from the previous year. Year 2008 is a special year in my sample, because towards the end of the year the market indeed experienced a rare disaster with the three-month S&P500 return lower than $-1/3$. This turns out to be the only occurrence of a market crash throughout the entire sample period according to my definition in Section 6.1. Hence, for this year instead of focusing on the first quarter, I try to predict the state of each stock over the 91-day period from August 22, 2008 to November 20, 2008, during which the S&P500 return was -41% .

For each quarter-length period mentioned above, I determine the states of each individual stock and the S&P500 index (H or L) over the period based on their returns using the definition of Section 6.1. Then, for each stock and each quarter-length period, I compare two quantities: $\mathbf{1}_{\text{stock in state } L}$, a dummy variable that takes on a value of 1 if the stock delivers a disastrous return over the period and 0 otherwise, and $\hat{\theta}(L|\text{realized state of S\&P500})$, the conditional disaster probability of the stock estimated from the previous year given the realized state of S&P500 over the period. I first check the correlation between these two quantities and find that they exhibit a strong positive correlation of 0.52. This indicates that the higher the predicted conditional disaster probability, the higher the actual occurrence of a firm disaster.

I then compute the following difference

$$\Delta = \mathbf{1}_{\text{stock in state } L} - \hat{\theta}(L|\text{realized state of S\&P500}).$$

Since $\mathbf{1}_{\text{stock in state } L}$ is a dummy variable for whether a firm disaster occurs and $\hat{\theta}(L|\text{realized state of S\&P500})$ represents the conditional disaster probability, their difference Δ for any one occasion does not contain much information. However, when I take the average across different stocks and periods, according to the law of large numbers these two quantities should be close to each other if my estimations are meaningful. Motivated by this idea, I compute the

average of Δ for normal and disaster states of the market separately. Conditional on a realized disaster market state (August 22, 2008 to November 20, 2008) and a realized normal market state (all other quarter-length periods), the average values of Δ are -0.01 and 0.13, respectively. These numbers are reasonably small in magnitude, providing evidence for good out-of-sample validity.

Next I compute summary statistics for the disaster risk variables Dis^M , Dis and $SysDis$. The results are reported in Table 1. The mean value of the market disaster risk Dis^M across the sample period is around 0.013, implying that on average a market crash is believed to occur with a risk-neutral probability of 0.013. In addition, the standard deviation, median, minimum, and maximum of Dis^M are estimated as 0.010, 0.012, 0.002, and 0.033, respectively. For the unconditional disaster risk (Dis) and the systematic disaster risk ($SysDis$) of individual stocks, I compute the mean and standard deviation across all stocks for each single year, and I report summary statistics of these annual cross-sectional means and standard deviations in the table. In particular, the mean value of the annual cross-sectional mean of Dis is equal to 0.031, indicating that on average individual stocks are believed to experience a disastrous event with a risk-neutral probability of 0.031. The mean value of the annual cross-sectional mean of $SysDis$ is around 0.335, suggesting that the conditional disaster probability of an individual stock given a market disaster is on average higher by 0.335 than its conditional disaster probability given normal market conditions.

Table 2 reports the correlation coefficients of $SysDis$ with other stock characteristics.⁴ The results show that $SysDis$ has a strong positive correlation of 0.40 with the unconditional disaster risk (Dis), meaning that stocks with a higher sensitivity to the market disaster risk are also prone to rare disasters by themselves. In addition, $SysDis$ also has a positive correlation of 0.14 with the CAPM beta ($Beta$), indicating that stocks with a higher sensitivity to the market disaster risk are also more sensitive to the overall market movement. Further, $SysDis$ has a small positive correlation with the book-to-market ratio ($B2M$), and is negatively correlated with coskewness ($Coskew$) and illiquidity ($IlliqMA$).

I next perform the Fama-MacBeth test on the cross-sectional risk premium associated

⁴Since $SysDis$ is estimated annually, I examine its correlation with other annually estimated stock characteristics only.

with $SysDis$. The test period extends from January 1997 to December 2012, which covers 16 years (192 months). In each month, the number of stocks included in the cross-sectional regression varies from 553 to 2270. The test results for the entire sample period are reported in Table 3. I first conduct the test with $SysDis$ as the only regressor (column 1). The result shows that $SysDis$ has a positive coefficient, but it is statistically insignificant. In fact, it remain insignificant after controlling for other stock characteristics (column 2). Consistent with existing evidence, beta has no significant effect on expected stock returns, and its coefficient is even negative. Further, $Coskew$ has a negative coefficient, as expected, but it is also statistically insignificant. Perhaps due to the relatively short sample period and the fact that I focus on the set of stocks with option trading only, some well-documented risk factors including size, book-to-market, and illiquidity do not show a significant effect, either. The one-month lagged return, on the other hand, has a significantly negative effect on future expected stock returns, implying a reversal in the return process.

I then examine how the relation between systematic disaster risk and expected stock returns changes with the market disaster risk. For this purpose, I rank all months in the sample by the market disaster risk (Dis^M) of the previous year and divide the entire sample into four subsamples, each containing 48 months. I repeat the test for each of the four subsamples separately. The results show that for the first three subsamples for which the lagged market disaster risk is below the first quartile (Columns 1 and 2), between the first and the second quartiles (Columns 3 and 4), and between the second and the third quartiles (Columns 5 and 6), $SysDis$ does not carry a positive premium. In fact, when the lagged market disaster risk is between the first and the second quartiles (Columns 3 and 4), $SysDis$ even has a negative coefficient significant at the 10% level after controlling for other risk factors. However, if I restrict the test to months for which the previous year market disaster risk is higher than the third quartile (columns 7 and 8), $SysDis$ becomes significant at the 1% level, and this effect remains after including other risk factors. Indeed, in this case increasing $SysDis$ by one standard deviation raises the expected monthly stock return by 63 basis points, which is equivalent to an increase of over 7% per year. Interestingly, when restricted to this subsample, the unconditional disaster risk (Dis) also carries a significantly positive premium. This implies that investors also require compensation for holding stocks with high unconditional disaster risk when a market crash is considered likely.

To sum up, the empirical study in this section shows no clear evidence that systematic disaster risk is priced by investors under normal market conditions. However, it has a strongly positive effect on expected future stock returns when the market-wide disaster risk is considerably high.

7 Application II: Bank Systemic Exposure

This section applies my approach to study the systemic risk of the banking sector. Systemic risk is the risk of collapse of the entire banking system due to the interrelation and interdependence of all banks. The study of systemic risk has recently attracted much attention, especially following the financial crisis of 2007–2009. Two aspects of systemic risk are of particular interest to researchers, one looking at the contribution of individual banks to the risk of the whole banking system, and the other focusing on the exposure of individual banks to systemic shocks.

Despite the general interests in learning about systemic risk, the empirical estimation is challenging due to the small probability of disastrous events. Different methods have been proposed in the literature to tackle this problem. For example, Huang, Zhou, and Zhu (2009) measure systemic risk by the price of insurance against financial distress, in which the default correlation between banks is proxied by the equity return correlation. Acharya, Pedersen, Philippon and Richardson (2010) propose the systemic expected shortfall (*SES*) measure, which estimates the propensity of a bank to be undercapitalized when the system as a whole is undercapitalized. Their approach relies on the power law distribution to model the tails. In addition, Adrian and Brunnermeier (2011) define the $\Delta CoVaR$ measure as the difference between the value-at-risk of the banking system conditional on an individual bank being in distress and the value-at-risk of the banking system conditional on the bank being solvent. The empirical estimation of $\Delta CoVaR$ uses quantile regression to capture the tails. One commonality of these papers is that they all measure systemic risk based on historical equity returns.

In this section, I examine the exposure of banks to systemic shocks and ask what bank characteristics are related to banks' systemic exposure. To this end, I propose a “systemic exposure” measure based on the conditional return distribution of a bank given the return

of the banking sector portfolio, which can be estimated using my approach by taking the sector portfolio as the related index. Since the measure is empirically estimated using option prices as opposed to historical equity returns, it is forward-looking and naturally accounts for investors' beliefs on the likelihood of disastrous events even if such events do not occur within sample.

7.1 Measure of Systemic Exposure

The systemic exposure measure resembles the systematic disaster risk measure introduced in Section 6.1. The main difference is that I now focus on the joint performance of a bank and the whole banking system, and hence I use the banking sector portfolio (as opposed to the market portfolio) as the related index.

As before, I look at equity returns over periods of three months. For both the banking sector and individual banks, I define two states ($N = K = 2$), the normal state (H) and the disaster state (L). I choose the disaster thresholds for the three-month returns of both the sector portfolio and individual banks to be $-1/2$. Intuitively, the sector (a bank) experiences a disastrous event if it loses more than one half of its original value within three months. Then, for each bank I define the systemic exposure measure as

$$SysExp = \theta(L|L) - \theta(L|H), \quad (10)$$

where $\theta(L|L)$ and $\theta(L|H)$ stand for the disaster probability of the bank conditional on the sector being in the disaster and the normal states, respectively. This measure captures the extent to which a bank is more likely to experience a disaster when the whole sector falls from the normal state to the disaster state. Put differently, it is a measure of the sensitivity of individual bank returns to the sector-wide disaster risk.

The systemic exposure measure takes values from -1 to 1 . A positive value means that a bank is more likely to experience a disaster when the sector as a whole is in the disaster state relative to when the sector performs normally. The higher the value of $SysExp$, the more sensitive the bank is to sector-wide disaster risk. On the other hand, a negative value of $SysExp$ implies that the bank becomes safer when the sector is hit by a disastrous event.

7.2 Data and Estimation

I use the KBW Bank Index as a proxy for the banking sector portfolio, which consists of 24 banking companies and has options actively traded on the Philadelphia Stock Exchange. I focus on computing the systemic exposure measure for each of the 24 current constituent banks of the index.⁵ The current composition of the index is shown in Table 4.

My sample period is from January 1996 to December 2012. For both the KBW Bank Index and the 24 banks, I collect daily information on option prices, security prices and dividend distributions from the OptionMetrics database. Notice that some of the 24 banks were formed by mergers and acquisitions at some points during my sample, and some others temporarily stopped trading on the option market for some periods. As a result, not all of the 24 banks have full records throughout the entire sample period.

I estimate $SysExp$ in three steps. First, on each date I compute the option-implied risk-neutral probability of disaster for both the KBW Index and the individual banks. To maintain the accuracy of estimation, I compute the risk-neutral disaster probability for a bank on a date only when (1) there are at least ten different options with pricing information available for the bank on that date, (2) the lowest moneyness level of available options for the bank on that date is no higher than 0.95, and (3) the shortest (longest) time to maturity of available options for the bank on that date is lower (higher) than three months. The average risk-neutral disaster probabilities for the KBW Bank Index and for the individual banks throughout the sample period are estimated as 0.010 and 0.018, respectively. In the second step, I run the constrained linear regression (8) for each bank over time to obtain the conditional disaster probabilities of the bank given different states of the sector. Finally, I compute $SysExp$ using the definition (10).

I am also interested in what bank characteristics predict banks' systemic exposure. To examine this issue, I construct the following bank characteristics. The first three are computed using daily security pricing information taken from the OptionMetrics database, whereas the rest are constructed from quarterly bank fundamentals obtained from the Compustat database.

⁵While I focus on the 24 current constituent banks of the KBW index in this application, there is nothing preventing me from estimating $SysExp$ for other banks. The 24 banks are among the largest banks in the sector and are considered to have the most influences upon the economy. Hence, results based on these banks are of particular economic import.

- Market beta ($Beta$), estimated with respect to the market portfolio using daily equity returns over the most recent 90 business days. The market excess returns and the risk-free rates needed for the estimation are taken from the Kenneth French online data library.
- Sector beta ($Beta^S$), estimated with respect to the banking sector portfolio using daily equity returns over the most recent 90 business days. The banking sector portfolio is proxied by the KBW Bank Index.
- Equity return volatility (Vol), estimated over the most recent 90 business days as the standard deviation of the daily equity returns.
- Bank size ($Size$), computed as the logarithm of the total market capitalization.
- Leverage ($Lever$), estimated as the ratio of the total book value of assets to the total book value of equity.
- Market-to-book ratio ($M2B$), computed as the ratio of the total market capitalization to the total book value of equity.
- Non-interest to interest income ratio ($N2I$), estimated as the ratio of total non-interest income to the total interest income. The inclusion of this variable is motivated by Brunnermeier, Dong and Palia (2012), who show that the $N2I$ ratio contributes to banks' systemic risk.
- Maturity mismatch. I construct two proxies for maturity mismatch. First is the ratio of total demand deposits to the total book value of assets ($MatMis1$), and the second is the ratio of total interest-bearing deposits to the total book value of assets ($MatMis2$). Since interest-bearing deposits generally have longer maturities and are less often withdrawn than demand deposits, the extent of maturity mismatch of a bank is positively related to the first proxy and negatively related to the second.

7.3 Results

I first estimate the systemic exposure of each bank using the entire sample period. That is, I run the constrained linear regression (8) based on the risk-neutral disaster probabilities of

both the KBW Index and the bank of interest throughout the whole sample. Notice that due to missing option prices for some banks during some periods of the sample, the effective estimation period may differ from bank to bank. Table 5 reports the estimation results. A first observation is that the systemic exposure is positive for all banks, indicating that banks are generally more likely to experience a disastrous event when the banking sector as a whole is hit by a disaster relative to when the sector operates in the normal state. For example, the systemic exposure of JPMorgan Chase & Co. is 0.686, meaning that when the banking sector moves from the normal state to the disaster state, the conditional disaster probability of the bank increases by 0.686. In addition, it can be seen that the systemic exposure varies much across different banks, from a minimum of 0.116 for Commerce Bancshares Inc. to a maximum of 0.993 for Comerica Inc.

I next ask the question of what bank characteristics predict systemic exposure. To address this question, I estimate *SysExp* for each bank at the end of each month based on option prices of the most recent 90 business days. To improve accuracy, I perform estimation only if the risk-neutral disaster probability is available for the bank for at least 60 days in the 90-business-day window. This yields a panel for the *SysExp* measure across different banks over time.

To make sure that my estimation of *SysExp* generates reasonable results, I check the out-of-sample validity by trying to predict future states of banks using my conditional disaster probability estimation. As in the previous application, I take the first quarter of each year and try to predict the state of a bank over the quarter based on the realized state of the sector during the same quarter and the conditional disaster probabilities estimated from the most recent 90-business-day window. From late 2008 to early 2009, the banking sector was hit by a disastrous shock, during which the three-month KBW return dropped below $-1/2$. This turns out to be the only occurrence of a sector-wide disaster throughout the entire sample period according to my definition in Section 7.1. Hence, instead of looking at the first quarter of 2009, I try to predict the state of each bank over the 91-day period from November 5, 2008 to February 3, 2009, during which the return of the KBW Index was -56% .

For each quarter-length period mentioned above, I determine the states of each bank and the KBW Index (H or L) over the period based on their returns using the definition of

Section 7.1. Then, for each bank and each quarter-length period, I compare two quantities: $\mathbf{1}_{\text{bank in state } L}$, a dummy variable that takes on a value of 1 if the bank delivers a disastrous return over the period and 0 otherwise, and $\hat{\theta}(L|\text{realized state of KBW})$, the conditional disaster probability of the bank estimated from the most recent 90-business-day window given the realized state of KBW over the period. I first check the correlation between these two quantities and find that they exhibit a strong positive correlation of 0.75. This indicates that the higher the predicted conditional disaster probability, the higher the actual occurrence of a bank disaster.

I then compute the following difference

$$\Delta = \mathbf{1}_{\text{bank in state } L} - \hat{\theta}(L|\text{realized state of KBW}),$$

and take the average of Δ across all banks and across different periods. My results show that conditional on a realized disaster state of KBW (November 5, 2008 to February 3, 2009) and a realized normal state of KBW (all other quarter-length periods), the average values of Δ are -0.01 and 0.0006, respectively. This serves as evidence for good out-of-sample validity of my estimation.

I then regress *SysExp* on lagged values of the bank characteristics discussed in Section 7.2. Since some characteristics have low frequency (quarterly available), I take as regressors the most recent value of each bank characteristic computed prior to the 90-business-day window used to estimate *SysExp*. Table 6 reports the regression results. I first run the regression without any fixed effects (column 1). The results show that the total return volatility, the market-to-book ratio, and the bank size have positive and statistically significant coefficients. On the other hand, the market beta, the leverage ratio, the non-interest to interest income ratio, and the maturity mismatch proxies do not seem to be related to banks' systemic exposure.

I then rerun the regression controlling for the bank-fixed effects (column 2). The return volatility and the market-to-book ratio remain positive and significant. However, the size variable loses its significance, and now it even switches to a negative sign. This implies that the positive correlation between size and systemic exposure is completely absorbed by the bank-fixed effects. In addition, the non-interest to interest income ratio becomes significant at the 10% level. Its positive coefficient indicates that the systemic exposure of

a bank increases with its noncore activities (e.g., investment banking, venture capital, and trading activities) relative to the traditional banking businesses.

Next, I perform the regression with both bank- and year-fixed effects. The market beta now has a significantly positive coefficient, indicating that a bank's sensitivity to the overall market movement positively predicts its systemic exposure. The return volatility continues to have a positive and significant coefficient, reinforcing the result that systemic exposure is increasing in the total return volatility. The market-to-book ratio loses its significance after controlling for the year-fixed effects. The coefficient of the size variable remains negative, and it now becomes significant at the 10% level, suggesting weak evidence for decreasing systemic exposure with size. Interestingly, the non-interest to interest income ratio continues to have a positive and even more significant effect on systemic exposure. This confirms the result that a bank's exposure to systemic shocks is driven by its non-traditional businesses. Finally, notice that after controlling for the bank- and year-fixed effects, the R-squared almost doubles, from 0.109 to 0.214, meaning that a large portion of the variation in the systemic exposure can be explained by bank- and year-specific factors.

Finally, I replace the market beta by the sector beta, controlling for all other bank characteristics and fixed effects. The sector beta has a positive coefficient and is significant at the 1% level, implying that banks that are more sensitive to the overall sector movement also tend to be more sensitive to the sector disaster risk. Results concerning other characteristics remain mostly unchanged, but the size variable now becomes insignificant with a t -statistic of -1.45. Overall, the results suggest that the systemic exposure cannot be fully predicted by the linear relation between individual bank returns and the sector returns (as captured by the sector beta) alone. Other bank characteristics such as the total return volatility and the non-interest to interest income ratio are also important predictors.

To sum up, I find that a bank exhibits higher exposure to systemic shocks when it is more sensitive to the overall market and sector movement, when its equity return becomes more volatile, and when it puts a larger weight on non-traditional businesses. I also find weak evidence that bank size negatively predicts the exposure to systemic shocks. In addition, banks' systemic exposure is also considerably driven by bank- and time-specific effects.

8 Application III: Ross Recovery for Individual Assets

The distributions of security returns generally differ under the physical versus the risk-neutral measures to reflect the adjustment for risk aversion. Ross (1976) and Breeden and Litzenberger (1978) show that option prices reveal information on the risk-neutral distributions of security returns (see Section 4.1). Unfortunately, it is a priori not clear how the physical return distributions can be estimated. Ross (2014) proposes a novel idea called the recovery theorem that enables us to separate physical return distributions from risk adjustments based on the risk-neutral distributions alone. Subsequent research further explores extensions and alternative approaches to tackle this problem (e.g., Carr and Yu (2012), Huang and Shaliastovich (2013), Martin and Ross (2013), Walden (2013), Dubynskiy and Goldstein (2013), and Borovička, Hansen and Scheinkman (2014)). All of these papers primarily focus on the market portfolio, and the recovery results do not directly apply to individual assets.

In this section, I extend the recovery results from the market portfolio to individual assets. I show that this can be accomplished through the risk-neutral joint distribution of the asset return with the market return, which can be conveniently estimated using my approach by taking the market portfolio as the related index. While I base my discussion on the original Ross (2014) recovery theorem, it is worth noticing that my extension is not restricted to Ross's original result but rather can be applied to any recovery result developed for the market portfolio. I start by reviewing the Ross recovery theorem.

8.1 Ross Recovery Theorem

Consider a one-period economy with two time points 1 and 2. Suppose that there is an aggregate market portfolio whose value defines N states indexed by $\{1, 2, \dots, N\}$.⁶ For any $n_1, n_2 \in \{1, 2, \dots, N\}$, let $\pi(n_1, n_2)$ denote the physical transition probability of moving from state n_1 at time 1 to state n_2 at time 2. In addition, let $\phi(n_1, n_2)$ denote the price at time 1 of an asset that pays one dollar at time 2 if the market is in state n_1 at time 1 and in state n_2 at time 2. Notice that $\phi(n_1, n_2)$ is typically referred to as the state price, and it is equal to the risk-neutral transition probability from state n_1 to state n_2 discounted at

⁶In Ross's model, it is the value of the market portfolio, instead of the return, that determines the state of nature.

the risk-free rate. The recovery theorem works directly with the state prices rather than the corresponding risk-neutral probabilities.

The Ross recovery theorem is motivated by the following model. Assume that a representative agent exists, who has an additively time-separable von Neumann-Morgenstern utility function $U(\cdot)$ with $U'(\cdot) > 0$ and $U''(\cdot) < 0$. Also assume that the representative agent has initial wealth W and a subjective time discount factor $\delta \in (0, 1)$. The agent chooses the optimal consumptions at time 1 and 2 subject to the budget constraint. Assume that the optimal consumption level depends on the state of the market only, and is in particular time independent. Denote the optimal consumption in state n by $c(n)$. Then, at time 1 the agent solves the following utility optimization problem:

$$\begin{aligned} & \max_{\{c(1), \dots, c(N)\}} U(c(n_1)) + \delta \sum_{n_2=1}^N U(c(n_2)) \pi(n_1, n_2) \\ & s.t. \\ & c(n_1) + \sum_{n_2=1}^N c(n_2) \phi(n_1, n_2) = W. \end{aligned}$$

The solution to the above optimization problem is given by the first order condition

$$U'(c(n_1)) \phi(n_1, n_2) = \delta U'(c(n_2)) \pi(n_1, n_2).$$

Rewriting the first order condition in matrix form, we obtain

$$\mathbf{D} \cdot \mathbf{\Phi} = \delta \mathbf{\Pi} \cdot \mathbf{D},$$

where

$$\mathbf{D} = \begin{pmatrix} U'(c(1)) & 0 & \dots & 0 \\ 0 & U'(c(2)) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & U'(c(N)) \end{pmatrix},$$

$$\mathbf{\Phi} = \begin{pmatrix} \phi(1,1) & \dots & \phi(1,N) \\ \vdots & \ddots & \vdots \\ \phi(N,1) & \dots & \phi(N,N) \end{pmatrix},$$

and

$$\mathbf{\Pi} = \begin{pmatrix} \pi(1,1) & \dots & \pi(1,N) \\ \vdots & \ddots & \vdots \\ \pi(N,1) & \dots & \pi(N,N) \end{pmatrix}.$$

Since $U'(\cdot) > 0$, the matrix \mathbf{D} is invertible. Then, solving for $\mathbf{\Pi}$ yields

$$\mathbf{\Pi} = \frac{1}{\delta} \mathbf{D} \cdot \mathbf{\Phi} \cdot \mathbf{D}^{-1}. \quad (11)$$

Since $\mathbf{\Pi}$ is a transition probability matrix, we must have that starting from any state, the transition probabilities to all states sum up to 1, i.e.,

$$\mathbf{\Pi} \cdot \mathbf{1}_{N \times 1} = \mathbf{1}_{N \times 1}.$$

This condition and (11) together imply

$$\mathbf{\Pi} \cdot \mathbf{1}_{N \times 1} = \frac{1}{\delta} \mathbf{D} \cdot \mathbf{\Phi} \cdot \mathbf{D}^{-1} \cdot \mathbf{1}_{N \times 1} = \mathbf{1}_{N \times 1},$$

which can be rewritten as

$$\mathbf{\Phi} \cdot \boldsymbol{\eta} = \delta \cdot \boldsymbol{\eta},$$

where $\boldsymbol{\eta} = \mathbf{D}^{-1} \cdot \mathbf{1}_{N \times 1}$. In particular, $U'(\cdot) > 0$ implies $\boldsymbol{\eta} \gg 0$, where \gg means “strictly greater than” in each and every element. That is, $\boldsymbol{\eta}$ is a strictly positive eigenvector of $\mathbf{\Phi}$ with a corresponding positive eigenvalue δ .

The Perron-Frobenius theorem (see Meyer (2000)) shows that all nonnegative irreducible matrices have a unique positive eigenvalue associated with a unique positive eigenvector up to scaling. If we assume no arbitrage, then the matrix of state prices $\mathbf{\Phi}$ is non-negative, with zero elements if and only if the corresponding physical probabilities are zero. If $\mathbf{\Phi}$ is also irreducible, one can uniquely identify the subjective discount factor δ and the vector $\boldsymbol{\eta}$ (or equivalently the diagonal matrix \mathbf{D}) up to scaling. Finally, plugging δ and \mathbf{D} into (11) allows one to recover the physical transition matrix $\mathbf{\Pi}$. Notice that while the matrix \mathbf{D} has an undetermined scaling parameter, the same parameter in \mathbf{D} and \mathbf{D}^{-1} exactly cancels out, rendering the physical transition matrix $\mathbf{\Pi}$ uniquely identified.

To apply the recovery theorem empirically, an essential step is to estimate the state price matrix $\mathbf{\Phi}$ from data. As discussed in Section 4.1, we can extract the risk-neutral probabilities and equivalently the corresponding state prices from option prices. One main challenge, however, is that this method only allows us to obtain the state prices starting from the realized initial state at time 1. Without loss of generality, assume that the realized

state of the market is $n_1 = 1$ at time 1. Then, one can estimate from option prices $\phi(1, n_2)$ for all n_2 (the first row of Φ) but not $\phi(n_1, n_2)$ for $n_1 \neq 1$.

To resolve this problem, Ross (2014) comes up with a brilliant idea. Let

$$\phi^l = \left(\phi^l(1, 1), \phi^l(1, 2), \dots, \phi^l(1, N) \right)$$

denote the vector of the l -period state prices starting from state 1. Specifically, $\phi^l(1, n)$ equals the price of an asset that pays one dollar in l periods if the market is in state 1 today and in state n after l periods. Notice that ϕ^1 coincides with $(\phi(1, 1), \phi(1, 2), \dots, \phi(1, N))$ from the above one-period setting. Further assume that the state of the market portfolio follows a Markov process. Then for any $l = 1, 2, \dots, N - 1$, we have the following recursive forward equations

$$\phi^{l+1} = \phi^l \cdot \Phi. \tag{12}$$

This forms a linear system of the state price matrix Φ , where the ϕ^l 's can be estimated from option prices with different maturities. Then, solving this linear system yields Φ , which can then be used to recover the physical transition matrix Π .

The recovery theorem deals with the entire transition matrix, which in particular includes transition probabilities starting from the realized initial state as well as from all other hypothetical initial states. Often times, however, we are mostly interested in the transition probabilities starting from the realized initial state $(\pi(1, 1), \pi(1, 2), \dots, \pi(1, N))$. Having the initial state fixed, we then have a 1–1 mapping between the market return and the future market value (given by (6)). Using the notations of my framework in Section 3 and focusing on transition probabilities starting from the realized initial state, (11) can be readily rewritten as

$$p_{t,t+1}^I(n) = \frac{1}{\delta} e^{-r_f} q_{t,t+1}^I(n) \frac{U'(c(1))}{U'(c(n))}, \tag{13}$$

where index I is chosen as the market portfolio, $p_{t,t+1}^I(n)$ is the physical probability of moving from the current market state to state n in one period, $e^{-r_f} q_{t,t+1}^I(n)$ is the state price, and δ as well as $\frac{U'(c(1))}{U'(c(n))}$ are uniquely determined from the Ross recovery result.

8.2 Extension to Individual Assets

To see how the recovery result can be applied to individual assets, it is useful to consider (13) for the market portfolio. This formula shows that the physical distribution of the market return can be decomposed into three parts. The first part is the reciprocal of the subjective discount factor $\frac{1}{\delta}$, which is a constant and in particular does not depend on the future market state. The second component is the state price $e^{-rf} q_{t,t+1}^I(n)$, which can be estimated from option prices. Finally, the third part is the pricing kernel $\frac{U'(c(1))}{U'(c(n))}$, representing the adjustment for risk aversion in different states of the market.

One cannot directly apply (13) to recover the physical return distribution of an individual asset from the associated risk-neutral distribution. The reason is that the same value of the asset return can appear in different states of the market unless there is perfect correlation between the asset return and the market return. Given that the pricing kernel $\frac{U'(c(1))}{U'(c(n))}$ is a function of the market return (fixing the initial market value), there is in general no single $\frac{U'(c(1))}{U'(c(n))}$ that can be used for any particular value of the asset return.

Fortunately, this problem can be resolved through the risk-neutral joint return distribution of the asset with the market. In particular, one can express the risk-neutral marginal distribution of the asset return as the sum of its risk-neutral joint probabilities with the market return, i.e.,

$$q_{t,t+1}(k) = \sum_{n=1}^N q(\tilde{r}_{t,t+1} = r(k), \tilde{r}_{t,t+1}^I = r^I(n)), \quad \forall k.$$

Then, each of the risk-neutral joint probabilities $q(\tilde{r}_{t,t+1} = r(k), \tilde{r}_{t,t+1}^I = r^I(n))$ can be adjusted for risk aversion by the appropriate value of $\frac{U'(c(1))}{U'(c(n))}$, which allows one to recover the physical joint probabilities. Finally, summing up the physical joint probabilities across all states of the market yields the physical marginal return distribution of the asset. This is formalized in the following proposition.

Proposition 3 *At any time t , the physical distribution of the asset return over one period head can be recovered as*

$$p_{t,t+1}(k) = \frac{1}{\delta} e^{-rf} \sum_{n=1}^N q(\tilde{r}_{t,t+1} = r(k), \tilde{r}_{t,t+1}^I = r^I(n)) \frac{U'(c(1))}{U'(c(n))}, \quad \forall k, \quad (14)$$

where δ and $\frac{U'(c(1))}{U'(c(n))}$ are uniquely determined from the Ross recovery theorem.

It is clear from Proposition 3 that the key to recovering the physical return distribution of an individual asset is the risk-neutral joint return distribution of this asset with the market portfolio. We know from the probability theory that this risk-neutral joint distribution equals the product of the risk-neutral marginal return distribution of the market and the conditional return distribution of the asset given the market return, i.e.,

$$q(\tilde{r}_{t,t+1}^I = r^I(n), \tilde{r}_{t,t+1} = r(k)) = q_{t,t+1}^I(n) \theta(k|n), \quad \forall n, k. \quad (15)$$

While the former can be obtained from option prices, my approach allows me to estimate the latter by taking the market portfolio as the related index. Hence, my framework lends itself naturally to the extension of the Ross recovery theorem to individual assets.

9 Conclusion

In this paper I propose a novel approach to recovering the conditional return distribution of an individual asset given the return of an aggregate index by regressing the marginal return distribution of the asset on that of the index. The identifying assumption that underlies this approach states that the conditional return distribution of the asset given any value of the index return does not vary over time. Intuitively, this means that the time variation in the return distribution of the asset is solely driven by the time variation in the return distribution of the index. This assumption allows me to make use of the time series information on the marginal return distributions of both securities to pin-down the time-invariant conditional distribution of the asset return given the index return. Empirically, I show how this approach can be implemented using option prices. I also discuss a variety of applications of this approach to the cross-sectional test of equity risk premium associated with systematic disaster risk, to the estimation of banks' exposure to systemic shocks, and to the extension the Ross (2014) recovery theorem to individual assets.

The advantage of my approach is that it generates the entire conditional return distribution of the asset given the index return, thus accounting for high moment properties and potential nonlinear patterns of the joint behavior of the two securities. In addition, since

the estimation relies on option prices rather than historical equity returns, it is forward-looking and reflects the effects of rare events even if they do not truly occur within sample. Further, compared to the copula approach, which is widely used to back out the joint distribution of multiple random variables from the associated marginals, my approach does not rely on any parametric assumptions on the security returns, thus allowing for greater flexibility. Yet, unlike some non-parametric procedures which are computationally very expensive, my approach can be easily implemented by a constrained linear regression.

Future work can be conducted in different directions. On the methodology side, research needs to be done on extending the framework to estimate the joint behavior of potentially more than two arbitrary economic variables (besides the returns of an asset and an aggregate index). In terms of applications, this approach can be used to address other important questions related to risk management and firm cyclicity, etc.

Appendix A

Proof of Proposition 1: At time t , the risk-neutral conditional probability of $\tilde{r}_{t,t+1} = r(k)$ given $\tilde{r}_{t,t+1}^I = r^I(n)$ is

$$q(\tilde{r}_{t,t+1} = r(k) | \tilde{r}_{t,t+1}^I = r^I(n)) = \frac{q(\tilde{r}_{t,t+1} = r(k), \tilde{r}_{t,t+1}^I = r^I(n))}{q_{t,t+1}^I(n)}. \quad (16)$$

By Assumption 2, the variation in the index return is the only priced risk. Thus, there exists a stochastic discount factor, whose value related to future payoffs depends on the future value of the index return only. Let $(\xi_{t,t+1}(1), \xi_{t,t+1}(2), \dots, \xi_{t,t+1}(N))$ denote the stochastic discount factor evaluated at time t related to payoffs to be received in different states of the index return at time $t+1$. This allows one to convert risk-neutral probabilities to physical probabilities. In particular, we have

$$q(\tilde{r}_{t,t+1} = r(k), \tilde{r}_{t,t+1}^I = r^I(n)) = e^{rf} \xi_{t,t+1}(n) p(\tilde{r}_{t,t+1} = r(k), \tilde{r}_{t,t+1}^I = r^I(n)), \quad (17)$$

and

$$q_{t,t+1}^I(n) = e^{rf} \xi_{t,t+1}(n) p_{t,t+1}^I(n). \quad (18)$$

Plugging (17) and (18) into (16) yields

$$q(\tilde{r}_{t,t+1} = r(k) | \tilde{r}_{t,t+1}^I = r^I(n)) = \frac{p(\tilde{r}_{t,t+1} = r(k), \tilde{r}_{t,t+1}^I = r^I(n))}{p_{t,t+1}^I(n)} = \theta(k|n),$$

as claimed. ■

Proof of Proposition 2: This proof relies on Theorem 2 of Liew (1976), which speaks about the estimates from a inequality-constrained least squares regression. The theorem says that when the inequality constraints are not binding with respect to the true parameter values, the estimates from the original inequality-constrained regression are equal to the estimates from the same regression without the inequality constraints when the sample size becomes large enough.

To prove the consistency of $\hat{\theta}_T$, I consider two possible cases.

Case 1: Suppose that $0 < \theta(k|n) < 1$ for all n and k .

According to Theorem 2 of Liew (1976), when the sample size T is large enough, $\hat{\theta}_T$ is equal to the coefficient estimates obtained from the following least squares linear regression problem with the equality constraint only

$$\mathbf{Q} = \mathbf{Q}^I \cdot \boldsymbol{\theta} + \boldsymbol{\epsilon}, \tag{19}$$

s.t.

$$\boldsymbol{\theta} \cdot \mathbf{1}_{K \times 1} = \mathbf{1}_{N \times 1}.$$

The exact formula for the equality-constrained least squares estimator can be found in Amemiya (1985) page 21. Since the equality constraint holds with respect to the true parameter values, it is standard to prove that under Assumptions 3 to 5, the equality-constrained least squares estimator is consistent. (The proof is similar to the proof for the consistency of the OLS estimator, and hence I omit it here for brevity.) Since $\hat{\theta}_T$ is equal to a consistent estimator when the sample size is large enough, $\hat{\theta}_T$ itself is also consistent.

Case 2: Suppose that the inequality constraint $0 \leq \theta(k|n) \leq 1$ is binding at one end with respect to the true parameter values for some n and k . Without loss of generality, consider a simplest case in which $\theta(1|1) = 0$ and $0 < \theta(k|n) < 1$ for all other pairs of n and k . The proofs for other cases are parallel.

Let $\hat{\theta}_T(1|1)$ denote the estimate of $\theta(1|1)$ obtained from the least squares regression problem (8) corresponding to a sample with size T . The inequality constraint implies $\hat{\theta}_T(1|1) \geq 0$. Then, for each $T > 0$ I have the following two subcases.

Subcase 2.1: Suppose that $\hat{\theta}_T(1|1) > 0$. In this scenario, $\hat{\theta}_T$ is equal to the coefficient estimates obtained from the following constrained least squares regression

$$\begin{aligned} \mathbf{Q} &= \mathbf{Q}^I \cdot \boldsymbol{\theta} + \boldsymbol{\epsilon}, \\ \text{s.t.} \\ \boldsymbol{\theta} \cdot \mathbf{1}_{K \times 1} &= \mathbf{1}_{N \times 1}, \\ \boldsymbol{\theta} &\leq \mathbf{1}_{N \times K}, \\ \theta(k|n) &\geq 0, \forall n, k \text{ unless } n = k = 1. \end{aligned}$$

Notice that the inequality constraint $\theta(1|1) \geq 0$ can be dropped because it is not binding ($\hat{\theta}_T(1|1) > 0$) and because I have a convex problem (convex objective function and compact domain). Theorem 2 of Liew (1976) indicates that when T is large enough, $\hat{\theta}_T$ is equal to the equality-constrained least squares estimator obtained from (19), which converges to the true parameter values $\boldsymbol{\theta}$ when the sample size approaches infinity.

Subcase 2.2: Suppose that $\hat{\theta}_T(1, 1) = 0$. In this scenario, $\hat{\theta}_T$ is equal to the coefficient estimates obtained from the following constrained least squares regression

$$\begin{aligned} \mathbf{Q} &= \mathbf{Q}^I \cdot \boldsymbol{\theta} + \boldsymbol{\epsilon}, \\ \text{s.t.} \\ \boldsymbol{\theta} \cdot \mathbf{1}_{K \times 1} &= \mathbf{1}_{N \times 1}, \\ \theta(1|1) &= 0, \\ \boldsymbol{\theta} &\leq \mathbf{1}_{N \times K}, \\ \theta(k|n) &\geq 0, \forall n, k \text{ unless } n = k = 1. \end{aligned}$$

According to Theorem 2 of Liew (1976), when T is large enough, $\hat{\theta}_T$ is equal to the coefficient estimates obtained from conducting the following regression problem with equality constraints only

$$\begin{aligned} \mathbf{Q} &= \mathbf{Q}^I \cdot \boldsymbol{\theta} + \boldsymbol{\epsilon}, \\ \text{s.t.} \\ \boldsymbol{\theta} \cdot \mathbf{1}_{K \times 1} &= \mathbf{1}_{N \times 1}, \\ \theta(1|1) &= 0. \end{aligned}$$

As in Case 1, since the equality constraints hold with respect to the true parameter values, the estimator from this equality-constrained problem converges to θ when the sample size approaches infinity.

The analyses of the above two subcases show that in Case 2, the sequence of estimators $\{\hat{\theta}_T\}_T$ can be viewed as consisting of two subsequences (each corresponding to one of the two subcases), both of which converge to the same true parameter values θ as the sample size approaches infinity. Therefore, the sequence $\{\hat{\theta}_T\}_T$ as a whole must also converge to θ .

Cases 1 and 2 together complete the proof. ■

Appendix B

As mentioned in Section 4.1, two technical issues need to be dealt with for the empirical estimation of the asset return distribution $\mathbf{q}_{t,t+1}$.

First, the estimation of the risk-neutral CDF (5) relies on differentiating the option price with respect to the strike price. Unfortunately, we are generally not able to obtain a close-form expression for this derivative. To resolve this issue, I resort to a linear approximation based on two points X^- and X^+ around the strike price X of interest, i.e.,

$$F(X) \approx e^{r_f} \frac{Put(X^+) - Put(X^-)}{X^+ - X^-}.$$

For all empirical applications in the paper, I choose X^- and X^+ to be \$0.01 to the left and right of X .

A second empirical challenge has to do with obtaining European option prices. Nearly all individual stock options are American options. While indices are typically represented by European options, the market option prices are only available at discrete values of the strike price and time to maturity. Thus, reasonable estimates are needed for European option prices at any arbitrary point of interest. In the empirical applications, I adopt a simple and commonly used approach in the literature (e.g., Shimko (1993), Malz (1997), and Figlewski (2010)) of first fitting the implied volatility surface and then deriving the option prices using the fitted volatility. Specifically, I start with the implied volatility of all available options written on the asset on the date of estimation. Information on the

implied volatility is provided by the OptionMetrics database, which is computed by the BMS model for European options and by the Cox-Ross-Rubinstein (CRR) model (Cox, Ross and Rubinstein (1979)) for American options. I then fit the implied volatility surface for a range of strike prices and maturities of interest based on the implied volatility of available options. Finally, I plug the fitted volatility at any particular point back into the BMS pricing model to obtain an estimation of the European option price.

I fit the implied volatility surface by kernel smoothing using a procedure similar to that adopted by the OptionMetrics database. On each date of estimation, I collect market pricing information of all options written on the same underlying asset indexed by $h = 1, 2, \dots, H$. For each option h , let σ^h represent the implied volatility, and let V^h be the vega of the option (which measures the sensitivity of option price to the volatility). Denote by $mn^h = X^h/S_t$ the moneyness of the option, by mt^h the time to maturity in years, and by cp^h a dummy variable that equals 0 for call options and 1 for put options. Then at any point with moneyness mo^* (within the range of moneyness of available options), time to maturity ma^* (within the range of time to maturity of available options), and call-put indicator cp^* , the estimated volatility $\hat{\sigma}$ takes the form

$$\hat{\sigma}(mn^*, mt^*, cp^*) = \frac{\sum_{h=1}^H V^h \sigma^h \Psi(mn^* - mn^h, mt^* - mt^h, cp^* - cp^h)}{\sum_{h=1}^H V^h \Psi(mn^* - mn^h, mt^* - mt^h, cp^* - cp^h)}, \quad (20)$$

where the kernel function Ψ is given by

$$\Psi(x, y, z) = \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{x^2}{2c_1} - \frac{y^2}{2c_2} - \frac{z^2}{2c_3}\right].$$

I naively choose $c_1 = c_2 = c_3 = 0.001$. These parameter values have been checked to generate reasonable fitting.

The idea of the kernel smoothing procedure is intuitive. At any point of interest, I estimate volatility as the weighted average of the implied volatility of all observed options, where observations close to the point of interest in terms of moneyness, time to maturity, and the call-put indicator are assigned higher weights than observations far away. In addition, since I am eventually interested in estimating the option prices, I would also want to assign higher weights to observed options whose prices have higher sensitivity to the volatility. To this end, I also include the option vega into the weights.

The kernel smoothing formula (20) works for moneyness and time to maturity levels within the range of available options. In my empirical applications, I mostly focus on rare disasters. To estimate the risk-neutral probabilities of rare disasters, I need European option prices (and hence the volatility) around the disaster thresholds, which feature very low moneyness levels, sometime even outside the range of moneyness for available options. When the moneyness at the disaster threshold is lower than the lowest observed moneyness of available options, I estimate volatility at the lowest observed moneyness level by (20) and use it as a proxy for the volatility around the disaster threshold.

Finally, notice from (20) that to perform kernel smoothing, besides the moneyness and time to maturity of interest, I also need to specify the call-put indicator cp^* . Theoretically, the implied volatility should be identical for a call option and a put option with the same moneyness and time to maturity. However, this is rarely the case in reality. A priori, it is not clear whether I should use the call or the put value for the estimation. In practice, most options trade at-the-money (ATM) or out-of-the-money (OTM), rendering the prices more accurate for ATM and OTM options. I thus adopt the following rule of thumb. When the moneyness of interest is lower than one ($mn^* < 1$), I estimate volatility for the put option ($cp^* = 1$), which trades OTM. When the moneyness of interest is higher than one ($mn^* > 1$), I estimate volatility for the call option ($cp^* = 0$), which again trades OTM. When the moneyness of interest is equal to one ($mn^* = 1$), both the call and the put options are exactly ATM. I then estimate volatilities for both the call and the put options and take the average.

References

- [1] Acharya, Viral V., Lasse H. Pedersen, Thomas Philippon, and Mathew Richardson, 2010. Measuring systemic risk. Working Paper, NYU Stern School of Business.
- [2] Adrian, Tobias, and Markus K. Brunnermeier, 2011. CoVaR. Fed Reserve Bank of New York Staff Reports.
- [3] Ait-Sahalia, Yacine, and Andrew W. Lo, 1998. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance* 53, 499–547.

- [4] Amemiya, Takeshi, 1985. *Advanced Econometrics*. Harvard University Press.
- [5] Amihud, Yakov, 2002. Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets* 5, 31–56.
- [6] Ang, Andrew, and Joseph Chen, 2002. Asymmetric correlations of equity portfolios. *Journal of Financial Economics* 63, 443–494.
- [7] Bali, Turan G., Ozgur Demirtas, and Haim Levy, 2009. Is there an intertemporal relation between downside risk and expected returns? *Journal of Financial and Quantitative Analysis* 44, 883–909.
- [8] Barro, Robert J., 2006. Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics* 121, 823–866.
- [9] Barro, Robert J., 2009. Rare disasters, asset prices, and welfare costs. *American Economic Review* 99, 243–264.
- [10] Black, Fischer, and Myron Scholes, 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 637–654.
- [11] Breeden, Douglas T., and Robert H. Litzenberger, 1978. Prices of state-contingent claims implicit in option prices. *Journal of Business* 51, 621–651.
- [12] Brunnermeier, Markus K., Gang Dong, and Darius Palia, 2012. Banks’ non-interest income and systemic risk. Working Paper, Princeton University and Rutgers Business School.
- [13] Bollerslev, Tim, and Viktor Todorov, 2011. Tails, fears, and risk premia. *Journal of Finance* 66, 2165–2211.
- [14] Borovička, Jaroslav, Lars Peter Hansen, and Jose A. Scheinkman, 2014. Misspecified recovery. Working Paper, New York University, University of Chicago, Columbia University and Princeton University.
- [15] Carr, Peter, and Jiming Yu, 2012. Risk, Return, and Ross Recovery. *Journal of Derivatives* 20, 38–59.

- [16] Chen, Hui, Scott Joslin, and Ngoc-Khanh Tran, 2012. Rare disasters and risk sharing with heterogeneous beliefs. *Review of Financial Studies* 25, 2189–2224.
- [17] Cox, John C., Stephen A. Ross, and Mark Rubinstein, 1979. Option pricing: a simplified approach. *Journal of Financial Economics* 7, 229–263.
- [18] Driessen, Joost, Pascal Maenhout, and Grigory Vilkov, 2013. Option-implied correlations and the price of correlation risk. Working Paper, Tilburg University, INSEAD, and Goethe University Frankfurt.
- [19] Dubynskiy, Sergey, and Robert S. Goldstein, 2013. Recovering drifts and preference parameters from financial derivatives. Working Paper, University of Minnesota.
- [20] Fama, Eugene F., and Kenneth R. French, 1992. The cross-section of expected stock returns. *Journal of Finance* 47, 427–465.
- [21] Fama, Eugene F., and Kenneth R. French, 1993. Common risk factors in the returns on bonds and stocks. *Journal of Financial Economics* 33, 3–56.
- [22] Fama, Eugene F., and James D. MacBeth, 1973. Risk, return, and equilibrium: empirical tests. *Journal of Political Economy* 81, 607–636.
- [23] Figlewski, Stephen, 2010. Estimating the implied risk-neutral density of the U.S. market portfolio. *Volatility and Time Series Econometrics*. Oxford University Press.
- [24] Gabaix, Xavier, 2008. Variable rare disasters: a tractable theory of ten puzzles in macro-finance. *American Economic Review* 98, 64–67.
- [25] Gabaix, Xavier, 2012. Variable rare disasters: an exactly solved framework for ten puzzles in macro-finance. *Quarterly Journal of Economics* 127, 645–700.
- [26] Gourio, François, 2012. Disaster risk and business cycles. *American Economic Review* 102, 2734–2766.
- [27] Harvey, Campbell R., and Akhtar Siddique, 2000. Conditional skewness in asset pricing tests. *Journal of Finance* 55, 1263–1295.

- [28] Hong, Yongmiao, Jun Tu, and Guofu Zhou, 2007. Asymmetries in stock returns: statistical tests and economic evaluation. *Review of Financial Studies* 20, 1547–1581.
- [29] Huang, Darien, and Ivan Shaliastovich, 2013. Risk adjustment and the temporal resolution of uncertainty: evidence from options markets. Working Paper, University of Pennsylvania.
- [30] Huang, Xin, Hao Zhou, and Haibin Zhu, 2009. A framework for assessing the systemic risk of major financial institutions. *Journal of Banking and Finance* 33, 2036–2049.
- [31] Jackwerth, Jens Carsten, and Mark Rubinstein, 1996. Recovering probability distributions from option prices. *Journal of Finance* 51, 1611–1631.
- [32] Jackwerth, Jens Carsten, 1999. Option-implied risk-neutral distributions and implied binomial trees: a literature review. *Journal of Derivatives* 7, 66–82.
- [33] Jegadeesh, Narasimham, and Sheridan Titman, 1993. Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance* 48, 65–91.
- [34] Jiang, Lei, Ke Wu, and Guofu Zhou, 2014. Asymmetry in stock returns: an entropy measure. Working Paper, Tsinghua University, Emory University, and Washington University in St. Louis.
- [35] Jorion, Philippe, 2000. Risk management lessons from long-term capital management. *European Financial Management* 6, 277–300.
- [36] Kadan, Ohad, Fang Liu, and Suying Liu, 2014. Generalized systematic risk. Working Paper, Washington University in St. Louis.
- [37] Kelly, Bryan, and Hao Jiang, 2013. Tail risk and asset prices. NBER Working Paper No. 19375.
- [38] Koenker, Roger, 2005. *Quantile Regression*. Cambridge University Press.
- [39] Kraus, Alan, and Robert H. Litzenberger, 1976. Skewness preference and the valuation of risk assets. *Journal of Finance* 31, 1085–1100.

- [40] Liew, Chong Kiew, 1976. Inequality constrained least-squares estimation. *Journal of the American Statistical Association* 71, 746–751.
- [41] Lintner, John, 1965a. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics* 47, 13–37.
- [42] Lintner, John, 1965b. Security prices, risk, and maximal gains from diversification. *Journal of Finance* 20, 587–615.
- [43] Longin, François, and Bruno Solnik, 2001. Extreme correlation of international equity markets. *Journal of Finance* 56, 649–676.
- [44] Malz, Allan M., 1997. Estimating the probability distribution of the future exchange rate from option prices. *Journal of Derivatives* 5, 18–36.
- [45] Martin, Ian, and Stephen A. Ross, 2013. The long bond. Working Paper, Stanford University and Massachusetts Institute of Technology.
- [46] Melick, William R., and Charles P. Thomas, 1997. Recovering an asset’s implied PDF from option prices: an application to crude oil during the Gulf crisis. *Journal of Financial and Quantitative Analysis* 32, 91–115.
- [47] Merton, Robert C., 1973. Theory of rational option pricing. *Bell Journal of Economics and Management Science* 4, 141–183.
- [48] Meyer, Carl D., 2000. *Matrix Analysis and Applied Linear Algebra*. SIAM.
- [49] Mossin, Jan, 1966. Equilibrium in a capital asset market. *Econometrica* 34, 768–783.
- [50] Nelsen, Roger B., 1999. *An introduction to copulas*. Springer, New York.
- [51] Posner, Steven E., and Moshe A. Milevsky, 1998. Valuing exotic options by approximating the SPD with higher moments. *Journal of Financial Engineering* 7, 109–125.
- [52] Roll, Richard, 1988. The international crash of October 1987. *Financial Analysts Journal* 44, 19–35.

- [53] Ross, Stephen A., 1976. Options and efficiency. *Quarterly Journal of Economics* 90, 75–89.
- [54] Ross, Stephen A., 2014. The recovery theorem. *Journal of Finance*, forthcoming.
- [55] Rubinstein, Mark, 1973. The fundamental theorem of parameter-preference security valuation. *Journal of Financial and Quantitative Analysis* 8, 61–69.
- [56] Rubinstein, Mark, 1998. Edgeworth binomial trees. *Journal of Derivatives* 1998, 20–27.
- [57] Ruenzi, Stefan, and Florian Weigert, 2013. Crash sensitivity and the cross-section of expected stock returns. Working Paper, University of Mannheim.
- [58] Sharpe, William F., 1964. Capital asset prices: a theory of market equilibrium under conditions of risk. *Journal of Finance* 19, 425–442.
- [59] Shimko, David, 1993. Bounds on probability. *Risk* 6, 33–37.
- [60] Siriwardane, Emil, 2013. The probability of rare disasters: estimation and implications. Working Paper, NYU Stern School of Business.
- [61] Skintzi, Vasiliki D., and Apostolos-Paul N. Refenes, 2005. Implied correlation index: a new measure of diversification. *Journal of Futures Markets* 25, 171–197.
- [62] Van Oordt, Maarten R.C., and Chen Zhou, 2012. Systematic tail risk. Working Paper, De Nederlandsche Bank and Erasmus University Rotterdam.
- [63] Wachter, Jessica A., 2013. Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance* 68, 987–1035.
- [64] Walden, Johan, 2013. Recovery with unbounded diffusion processes. Working Paper, University of California at Berkeley.

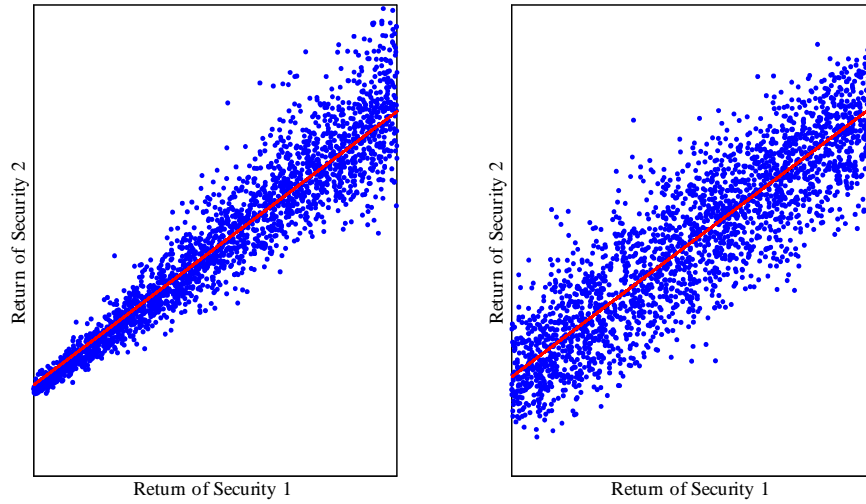


Figure 1: Second Moment Properties in Joint Return Behavior

This figure plots the returns of two pairs of hypothetical securities. Both pairs predict the same conditional mean return of one security given the return of the other. However, the second-moment patterns of the two pairs are distinct in the sense that the first pair has increasing correlation when the returns become lower, whereas the second pair has symmetric correlation in the entire region of returns.

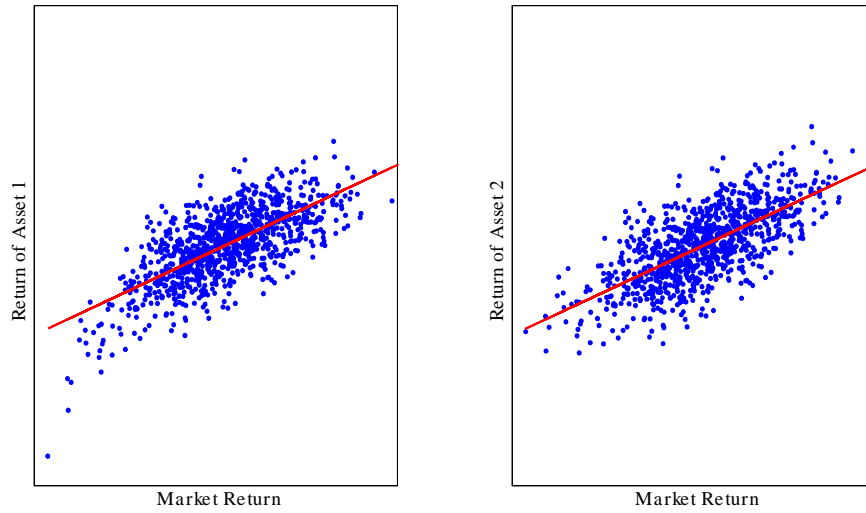


Figure 2: Nonlinearity in Joint Return Behavior

This figure plots the returns of two hypothetical assets against the market return. The returns of both assets fit the same linear relation with the market return. However, asset 1 is more sensitive to the market disaster risk than asset 2 in the sense that it tends to deliver lower returns when the market return becomes extremely low.

Table 1: Summary Statistics of Disaster Risk Variables

This table reports summary statistics of the disaster risk variables for the time period from 1996 to 2011. For the systematic disaster risk ($SysDis$) and the unconditional disaster risk (Dis) of individual assets, I compute in each year the cross-sectional mean and standard deviation of each variable, and I report in the table the mean, median, minimum, and maximum of the annual cross-sectional mean as well as the mean of the annual cross-sectional standard deviation. For the market disaster risk (Dis^M), the table reports the mean, standard deviation, median, minimum, and maximum of the variable over the sample period.

Variable	Mean of annual mean	Mean of annual S.D.	Median of annual mean	Min of annual mean	Max of annual mean
$SysDis$	0.3346	0.3643	0.2886	0.1164	0.6910
Dis	0.0314	0.0377	0.0295	0.0112	0.0577
Variable	Mean	S.D.	Median	Min	Max
Dis^M	0.0131	0.0100	0.0123	0.0015	0.0329

Table 2: Pairwise Correlations of Stock Characteristics

This table reports pairwise correlation coefficients of various stock characteristics for the time period from 1996 to 2011. The stock characteristics include the systematic disaster risk (*SysDis*), unconditional disaster risk (*Dis*), CAPM beta (*Beta*), coskewness (*Coskew*), book-to-market ratio (*B2M*), and illiquidity (*IlliQMA*).

	<i>SysDis</i>	<i>Dis</i>	<i>Beta</i>	<i>Coskew</i>	<i>B2M</i>	<i>IlliQMA</i>
<i>SysDis</i>	1					
<i>Dis</i>	0.4048	1				
<i>Beta</i>	0.1353	0.1104	1			
<i>Coskew</i>	-0.0604	-0.0323	0.0150	1		
<i>B2M</i>	0.0123	0.0248	0.0050	-0.0112	1	
<i>IlliQMA</i>	-0.0236	0.1044	0.0227	-0.0210	0.0251	1

Table 3: Systematic Disaster Risk Premium for Entire Sample

This table reports the Fama-MacBeth test results of the cross-sectional risk premium associated with systematic disaster risk ($SysDis_{t-1}$) for the time period from 1997 to 2012. The test controls for a variety of risk factors, including the unconditional disaster risk (Dis_{t-1}), CAPM beta ($Beta_{t-1}$), coskewness ($Coskew_{t-1}$), size ($Size_{t-1}$), book-to-market ratio ($B2M_{t-1}$), lagged return (Ret_{t-1}), and illiquidity ($IlliqMA_{t-1}$). The t -statistics are reported in the parentheses below the corresponding estimates. Asterisks denote statistical significance at the 1% (***) , 5% (**) and 10% (*) levels.

Variable	(1)	(2)
$SysDis_{t-1}$	0.0040 (1.06)	0.0017 (0.84)
Dis_{t-1}		0.0261 (0.46)
$Beta_{t-1}$		-0.0107 (-1.41)
$Coskew_{t-1}$		-2104.67 (-0.81)
$Size_{t-1}$		-0.0011 (-1.16)
$B2M_{t-1}$		-0.0001 (-0.80)
Ret_{t-1}		-0.0222 (-2.60)**
$IlliqMA_{t-1}$		0.0308 (1.17)
$Constant$	0.0063 (1.81)*	0.0333 (1.73)*

Table 4: Systematic Disaster Risk Premium and Market Disaster Risk

This table reports the Fama-MacBeth test results of the cross-sectional risk premium associated with systematic disaster risk ($SysDis_{t-1}$) for subsamples constructed based on the market disaster risk estimated from the previous year (Dis_{t-1}^M). Columns 1 and 2 report results for the subsample with Dis_{t-1}^M below the first quartile, columns 3 and 4 correspond to the subsample with Dis_{t-1}^M between the first and the second quartiles, columns 5 and 6 report results for the subsample with Dis_{t-1}^M between the second and the third quartiles, and columns 7 and 8 restrict the test to the subsample with Dis_{t-1}^M above the third quartile. The test controls for a variety of risk factors, including the unconditional disaster risk (Dis_{t-1}), CAPM beta ($Beta_{t-1}$), coskewness ($Coskew_{t-1}$), size ($Size_{t-1}$), book-to-market ratio ($B2M_{t-1}$), lagged return (Ret_{t-1}), and illiquidity ($IlliqMA_{t-1}$). The t -statistics are reported in the parentheses below the corresponding estimates. Asterisks denote statistical significance at the 1% (***) , 5% (**) and 10% (*) levels.

Variable	Low Dis_{t-1}^M		2		3		High Dis_{t-1}^M	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$SysDis_{t-1}$	-0.0012 (-0.26)	0.0013 (0.48)	-0.0042 (-0.53)	-0.0060 (-1.73)*	-0.0080 (-1.06)	-0.0056 (-1.59)	0.0293 (3.57)***	0.0172 (3.18)***
Dis_{t-1}		-0.0674 (-0.76)		0.0852 (0.68)		-0.1657 (-1.23)		0.2523 (2.77)***
$Beta_{t-1}$		-0.0080 (-1.16)		-0.0113 (-0.61)		-0.0139 (-0.90)		-0.0098 (-0.55)
$Coskew_{t-1}$		-1.0652.98 (-1.12)		1321.01 (0.37)		-26.07 (-0.01)		939.38 (0.92)
$Size_{t-1}$		0.0009 (0.76)		-0.0013 (-0.67)		-0.0007 (-0.27)		-0.0034 (-1.78)*
$B2M_{t-1}$		0.0001 (0.59)		-0.0002 (-0.58)		-0.0004 (-2.56)**		0.0001 (0.55)
Ret_{t-1}		-0.0052 (-0.38)		-0.0486 (-2.30)**		-0.0137 (-0.72)		-0.0213 (-1.62)
$IlliqMA_{t-1}$		0.0628 (2.43)**		0.0449 (0.84)		0.0367 (0.71)		-0.0210 (-0.29)
$Constant$	0.0085 (1.73)*	0.0030 (0.13)	-0.0004 (-0.04)	0.0305 (0.77)	0.0024 (0.33)	0.0272 (0.65)	0.0148 (2.55)**	0.0728 (1.57)

Table 5: Systemic Exposure of Banks

This table reports the systemic exposure ($SysExp$) of the 24 current constituent banks of the KBW Bank Index computed for the time period from 1996 to 2012. Due to missing option prices for some banks during some periods of my sample, the effective estimation period may differ from bank to bank.

Bank Name	$SysExp$
BB&T Corporation	0.7636
Bank of America Corp	0.9881
Capital One Financial Corp	0.9789
JPMorgan Chase & Co	0.6860
Citigroup	0.9880
Comerica Inc	0.9925
Commerce Bancshares Inc	0.1164
Cullen/Frost Bankers Inc	0.2079
Fifth Third Bancorp	0.9865
First Niagara Financial Group	0.2534
U.S. Bancorp	0.8038
Huntington Bancshares Inc	0.9907
KeyCorp	0.9883
M&T Bank Corporation	0.7331
Bank of New York Mellon	0.7632
New York Community Bank	0.5715
Northern Trust Corp	0.4948
PNC Financial Services Group	0.7805
People's United Financial Inc	0.1553
Regions Financial Corporation	0.9852
State Street Corporation	0.8056
SunTrust Banks Inc	0.9919
Wells Fargo & Company	0.9912
Zions Bancorporation	0.9831

Table 6: Systemic Exposure and Bank Characteristics

This table reports the coefficient estimates from regressing the systemic exposure (*SysExp*) of the 24 current constituent banks of the KBW Bank Index on lagged values of various bank characteristics for the time period from 1996 to 2012. The systemic exposure is computed for each bank at the end of each month using option prices of the most recent 90 business days. The bank characteristics include market beta (*Beta*), sector beta (*Beta^S*) equity return volatility (*Vol*), leverage ratio (*Lever*), market-to-book ratio (*M2B*), log market capitalization (*Size*), non-interest to interest income ratio (*N2I*), ratio of total demand deposits to total book value of assets (*MatMis1*), and ratio of total interest-bearing deposits to total book value of assets (*MatMis2*). Column 1 reports the regression results without fixed effects, column 2 controls for the bank-fixed effects, and column 3 controls for both bank- and year-fixed effects. The *t*-statistics are reported in the parentheses below the corresponding estimates. Asterisks denote statistical significance at the 1% (***) , 5% (**) and 10% (*) levels.

Variable	(1)	(2)	(3)	(4)
<i>Beta</i> _{<i>t</i>-1}	-0.0023 (-0.07)	0.0112 (0.32)	0.0587 (2.54)**	
<i>Beta^S</i> _{<i>t</i>-1}				0.2196 (4.35)***
<i>Vol</i> _{<i>t</i>-1}	6.5414 (6.20)***	5.4125 (6.10)***	4.1575 (2.97)***	3.1228 (2.29)**
<i>Lever</i> _{<i>t</i>-1}	0.0054 (0.93)	-0.0003 (-0.04)	-0.0077 (-0.89)	-0.0088 (-1.12)
<i>M2B</i> _{<i>t</i>-1}	0.0395 (2.99)***	0.0547 (3.55)***	0.0258 (0.81)	0.0270 (0.90)
<i>Size</i> _{<i>t</i>-1}	0.0349 (2.25)**	-0.0581 (-1.73)	-0.0739 (-1.81)*	-0.0574 (-1.45)
<i>N2I</i> _{<i>t</i>-1}	0.0199 (0.79)	0.0480 (1.84)*	0.0974 (3.02)***	0.1019 (3.29)***
<i>MatMis1</i> _{<i>t</i>-1}	-0.1616 (-0.71)	-0.3353 (-1.10)	0.5214 (1.25)	0.4235 (1.13)
<i>MatMis2</i> _{<i>t</i>-1}	-0.2059 (-0.88)	0.1767 (0.60)	0.0183 (0.07)	0.0421 (0.16)
<i>Constant</i>	-0.4852 (-1.16)	0.9277 (1.47)	1.1158 (1.54)	0.7549 (1.07)
Bank-fixed effects	N	Y	Y	Y
Year-fixed effects	N	N	Y	Y
Number of observations	3123	3123	3123	3123
R-squared	0.1092	0.1371	0.2141	0.2205