ORIGINAL PAPER

DOI: 10.26794/2587-5671-2019-23-1-38-48
УДК 339.7(045)=111
JEL G11, G24, C14

Can Stock Analysts Predict Market Risk? New Evidence From Copula Theory

I.S. Medovikov
Brock University, Canada, Ontario
https://orcid.org/0000-0002-9322-0410

ABSTRACT

We assess investment value of stock recommendations from the standpoint of market risk. We match I/B/E/S (Institutional Brokers' Estimates System) consensus recommendations issued in January 2015 for a cross-section of U.S. public equities with realized volatility of these papers, showing that these recommendations significantly correlate with subsequent changes in market risk. Thus, the results indicate that to some extent the analysts can predict an increase or decrease in risk, which can benefit asset management. However, the relationship between the recommendations and the risk is not linear and depends on the specific recommendation. Using a semi-parametric copula model, we find recommendation levels to be associated with future changes in volatility. We further find this relationship to be asymmetric and most pronounced among the best-rated stocks which experience largest volatility declines. We conduct a trading simulation showing how stock selection based on such ratings can lead to a reduction in portfolio-level value-at-risk.

Keywords: stock analysts; copulas; investment management; portfolio management; semi-parametric analysis

For citation: Medovikov I.S. Can stock analysts predict market risk? New evidence from copula theory. *Finansy: teoriya i praktika = Finance: Theory and Practice*. 2019;23(1):38-48. DOI: 10.26794/2587-5671-2019-23-1-38-48

ОРИГИНАЛЬНАЯ СТАТЬЯ

Могут ли фондовые аналитики предсказать рыночный риск? Новые сведения из теории копулы

И.С. Медовиков
Университет Брок, Онтарио, Канада
https://orcid.org/0000-0002-9322-0410

АННОТАЦИЯ

Статья оценивает способность финансовых аналитиков прогнозировать рыночный риск. Сопоставляя консенсус-рекомендации, выпущенные аналитиками для акций публичных компаний США, содержащихся в системе I/B/E/S (Institutional Brokers' Estimates System) на январь 2015 г., с фактической волатильностью этих бумаг, мы показываем, что эти рекомендации значимо коррелируют с последующими изменениями в уровне рыночного риска. Таким образом, наши результаты указывают на то, что аналитики хотя бы в какой-то степени способны предсказать нарастание или убывание риска, что может принести пользу в управлении активами. Однако взаимоотношение между рекомендациями и риском не является линейным и зависит от конкретной рекомендации. Используя семи-параметрическую статистическую модель на основе теории копул, автор показывает, что «экстремальные» рекомендации (т.е. самые положительные или самые отрицательные) несут гораздо большую информационную нагрузку, чем остальные. В контексте научной литературы на данную тему результаты исследования, по-видимому, представляют собой одну из первых попыток установить эмпирическую зависимость между рекомендациями аналитиков и рыночным риском.

Ключевые слова: фондовые аналитики; копула; управление инвестициями; управление портфелем; семи-параметрический анализ

Для цитирования: Medovikov I.S. Can stock analysts predict market risk? New evidence from copula theory. *Финансы: теория и практика*. 2019;23(1):38-48. DOI: 10.26794/2587-5671-2019-23-1-38-48

1. INTRODUCTION

While the stock recommendations and ratings issued by sell-side security analysts at major banks and brokerage houses continue to receive considerable amount of attention in the media and among investors, their investment value remains uncertain despite much research during the past twenty years. A substantial body of literature on this subject suggests that at least in some circumstances, following analyst recommendations may lead to superior portfolio returns (see, for example, [1–4], among others), but achieving this in practice may be less than straightforward as recommendation profitability seems to depend on multiple factors such as timely access to analysts' reports [5], speed of portfolio turnover [3], proximity of earnings announcements [6], recommendation revisions [7, 8], and also varies with recommendation levels [9].

To-date, this work has mostly focused on determining whether obtaining a rate of return in excess of some market benchmark is possible based on public stock ratings. But very little appears to be known about the value of ratings from the standpoint of market risk, which is surprising given the importance of risk management for any portfolio selection process.

The few studies that are available on this subject include [10] and more recently [11], who find the variance of stock returns to increase during periods of time surrounding recommendation revisions. While this is insightful, it is of limited use to practitioners seeking to use ratings to manage market risk on go-forward basis, as that would require establishing a link between recommendation levels and future levels of price volatility. At present, this link appears to remain unexplored and the aim of this paper is to fill this gap. In what follows we seek to answer several specific questions: does buying well-rated stocks lowers portfolio risk going forward? If so, at what horizon, and by how much? Is the relationship the same for all recommendations levels? We show how answers to these questions can aid portfolio stock selection from the risk management perspective.

To achieve this, we match I/B/E/S consensus recommendations issued for U.S.-listed companies during January 2015 with realized volatility of security returns up to one year following recommendations issue and use a flexible semiparametric conditional copula model to map the relationship between recommendation and subsequent changes in return volatility in high detail. We find recommendation levels to be associated with the change in volatility six to twelve months following recommendation issue, suggesting that consensus recommendations may indeed help manage portfolio risk. Further, we find this relationship to be more pronounced among best-rated securities which appear to experience largest volatility declines, while the converse does not seem to hold for worst-rated stocks. In all cases, in line with the earlier literature on the information content of analyst recommendations, we find this predictive ability to be conditional on recommendation changes, meaning that ratings representing a revision are better predictors of future risk.

Additionally, to assess investment value of our results in applied context we conduct a trading simulation where we model holding returns to two long-only equity portfolios, one consisting of the best-rated stocks, and an-other containing worst-rated securities. Both portfolios are held up to twelve months following recommendation issue and are not rebalanced. We estimate the realized 5% value-at-risk for both portfolios at several event horizons and find that in nearly-all cases the "best-rated" portfolio experienced a substantial reduction in value-at-risk relative to "worst-rated" holdings, providing a straightforward "recipe" for using I/B/E/S recommendations to lower port-folio exposure to market risk.

The remainder of this paper is organized as follows. Section 2 reviews empirical methodology adopted here and, specifically, details the construction and estimation of the copula model. Recommendations and security price data are described in Section 3. We present our empirical results in Section 4 and provide a discussion in Section 5.

2. METHODOLOGY

We follow an event-study approach and focus on a single release of the I/B/E/S consensus recommendations on January 15, 2015, which is the most recent vintage available at present that enables us to match recommendations with a full calendar year of subsequent security returns. I/B/E/S consensus recommendations are issued monthly and contain buy, sell, and hold-type ratings for a large number of public securities. These ratings are provided on a standardized numerical scale where smaller numbers represent more-favorable, and larger numbers represent less-favorable recommendations, and Section 3.1 provides a detailed review of I/B/E/S recommendations in general along with the summary statistics for the data in our sample.

For notational convenience, let $R_{i,t}$ denote the level of the consensus recommendation for security i observed on trading day t, and let $t = T$ denote recommendation issue date which in our case is January 15, 2015. For any two arbitrary trading days s, k such that $k > s$ and for any security i, let the realized variance of daily returns between these days as

$$v_i(s,k)=\sum_{j=s+1}^{k}\ln\left(p_{i,j}/p_{i,j-1}\right)^2, \qquad (1)$$

where $p_{i,j}$ is the adjusted closing price of *i*'th company stock on trading day *j*. While this represents a helpful metric from an investor's standpoint, when it comes to the assessment of predictive value of analyst recommendations, simple matching of such unconditional realized variance with recommendation levels may be misleading since there may be substantial differences in $v_i(s, k)$ among firms due to factors not fully reflected in recommendations such as, for example, industry focus of the firm. Consider a technology firm and a utilities firm that share the same, say, "strong buy' rating. Individual analysts issuing recommendations covering these firms tend focus on specific industries and rarely follow more than two or three sectors. They may therefore issue identical "strong buy" ratings if they see both firms as being likely to out-perform relative to their industry peers. But technology firms typically carry a higher level of market risk than utilities firms, which may therefore not be reflected in their recommendations. To address this, we standardize our volatility measures and instead use percent changes in realized variance during some window of time following recommendation issue. Given a set span of time consisting of *K* trading days, we define the change in realized volatility K days following recommendation issue denoted by $V_i(K)$ as a percentage difference

$$V_i(K)=100\begin{pmatrix} v_i\left(T^*+1,T^*+K\right)/ \\ /v_i\left(T^*-K,T^*-1\right)\end{pmatrix}-1, \qquad (2)$$

where $v_i(s, t)$ is defined as before. A value of $V_i(K) = 10$ therefore simply indicates that realized variance of *i*'th company stock returns in our sample increased by ten percent during *K* days following a new recommendation, compared to a period of the same length before. Note that this definition omits return realized on the day of recommendation issue. We are therefore omitting short-term price reaction that happens right on the day of recommendation release. Our aim in the following sections is to study the association between $R_{i,T}$ and $V_i(K)$, for various event horizons *K*.

2.1. The Copula Approach

To map the relationship between recommendations and changes in return volatility in high detail we adopt a modeling approach that is based on the statistical theory of copulas. Copulas have been widely adopted in finance and economics in recent years largely due to their ability to separately model marginal behavior of the variables from the interaction between them. In our case, this allows construction of a flexible distribution model for recommendations and volatility changes that can accommodate their very different marginal characteristics such as boundedness and degrees of skewness. While it would be otherwise difficult to find a known suitable bivariate distribution model, this is relatively straightforward to accomplish using the copula. As we show in Section 4, the ability to model the entire distribution in such a way provides a rich picture of dependence that can reveal asymmetries and possible nonlinearities in the recommendations-volatility relationship.

2.1.1. The Conditional Copula Model

Let $F(R_{i,T}; \theta_r)$ and $G(V_i(K); \theta_v)$ represent the marginal distribution functions of $R_{i,T}$ and $V_i(K)$ respectively and let $H(R_{i,T}, V_i(K); \theta_c)$ be their joint d.f., where is a vector containing all distribution parameters and θ_r and θ_v are parameter vectors for the marginals. Following [19], the function *H* can be represented as

$$H\left(R_{i,T^*},V_i(K);\theta\right)= \\ =C\left(F\left(R_{i,T^*};\theta_r\right),G\left(V_i(K);\theta_v\right);\theta_c\right), \qquad (3)$$

where *C* is the copula of *H*, and θ_c is a vector containing copula arameters. Letting $u = F(R_{i,T}; \theta_r)$ and $v = G(V_i(K); \theta_v)$ it should be evident that $C:[0,1]^2\to[0,1]$ is a joint distribution function of (u; v). The copula specifies how the marginals F and G are "coupled" together to form the joint d.f. H and as such provides a complete description of the dependence between $R_{i,T}$ and $V_i(K)$. The parameters in θ_c in turn capture the strength of association. Since by construction u and v are uniformly-distributed on [0; 1] this description is independent from the choice of the marginals and when the variables are continuous it is also unique. Decomposition in (3) can be further enriched by allowing the marginals to be conditional on some exogenous variables, turning C into a conditional copula. For further details on conditional copulas, see [13]. Our aim in this and the following sections is therefore to develop and estimate a suitable model of *C* for recommendations $R_{i,T}$ contained in the January 2015 I/B/E/S vintage and volatility changes $V_i(K)$.

Many common measures of association can be expressed in terms of the copula *C*. For example, rank correlation coefficients such as Kendall's and Spearman's can be written respectively as

$$\tau = 4\int_0^1\int_0^1 C(u,v)\,dC(u,v) - 1 \quad (4)$$

and

$$\rho = 12\int_0^1\int_0^1 C(u,v)\,du\,dv - 3. \quad (5)$$

Various families of copulas represent a variety of dependence structures. For example, the Gaussian copula captures symmetric linear correlation, while other common families such as Gumbel or Clayton copulas can capture dependence that is skewed toward upper or lower tail of the joint distribution. Copulas belonging to the Student family can capture linear correlation combined with symmetric tail dependence. For an excellent introduction to copulas see [14–17] provide an overview of applications of copulas to problems in finance.

2.2. Copula Maximum Likelihood Estimation

Once an appropriate parametric form for C is selected, parameters in θ_c can be estimated relatively easily using maximum likelihood. Note that differentiating (3) allows us to represent the joint PDF of R_{iT} and $V_i(K)$ as

$$h\left(R_{i,T^*}, V_i(K); \theta\right) = f\left(R_{i,T^*}; \theta_r\right) g\left(V_i(K); \theta_v\right) c\left(R_{i,T^*}, V_i(K); \theta_c\right), \quad (6)$$

where functions *f, g* are the marginal PDFs and *c* is the copula density defined as

$$c\left(R_{i,T^*}, V_i(K); \theta_c\right) = \frac{\delta C\left(F\left(R_{i,T^*}; \theta_r\right), G\left(V_i(K); \theta_v\right); \theta_c\right)}{\delta R_{i,T^*}\,\delta V_i(K)}. \quad (7)$$

Obtaining the estimates θr, θv, and θc therefore requires that we maximize the corresponding log-likelihood function

$$\log\left(h\left(R_{i,T^*}, V_i(K); \theta\right)\right) = \log\left(f\left(R_{i,T^*}; \theta_r\right)\right) + \log\left(g\left(V_i(K); \theta_v\right)\right) + \log\left(c\left(R_{i,T^*}, V_i(K); \theta_c\right)\right). \quad (8)$$

Operating on the entire log-likelihood in (9), while possible, can be computationally costly, and [18] propose an alternative two-step procedure where marginal models are estimated first using maximum likelihood, and the copula log-likelihood $\log\left(h\left(R_{i,T^*}, V_i(K); \theta\right)\right)$ is maximized second using first-step MLE estimates of the marginals $F(R_{i,T}; \theta_r)$ and $G(V_i(K); \theta_v)$. When the marginals are parametric, this procedure is known as Inference Functions for Margins (IFM), and with non-parametric margins the procedure is termed Canonical Maximum Likelihood (CML).

3. RECOMMENDATIONS AND PRICE DATA

3.1. I/B/E/S Consensus Recommendations

Our recommendations dataset consists of the entire I/B/E/S consensus file issued on January 15, 2015. I/B/E/S recommendations are released on the third Thursday of every month and in the case of January 15, 2015 vintage contain analyst recommendations and several related statistics for a total of 4,580 U.S. public companies. For every company in the file, every reporting analyst provides a rating which is mapped onto a standardized 1–5 scale, where smaller numbers correspond to more-favorable recommendations. A rating coded as 1 therefore represents the most favorable recommendation which is referred to as "strong buy", while 5 is the least favorable recommendation, or "strong sell". Intermediate recommendation levels include a neutral "hold" rating coded as 3 and weaker "buy" and "sell" ratings coded as 2 and 4 respectively.

Most companies are tracked by multiple analysts, and consensus recommendations for such firms are simply averages of all individual ratings. This represents a difficulty when interpreting consensus recommendations since analysts' coverage is highly skewed towards larger listed rms, leaving a large number of smaller firms very thinly-covered. Many stocks, particularly those sold over the counter are being covered by only a single analyst and almost half of all firms in our sample are covered by less than five. *Table 1* shows distribution of analyst coverage across firms in the January 15, 2015 I/B/E/S vintage.

Interestingly, analysts appear to show strong consensus when it comes to thinly-covered rms. In our sample we find that for 64% of firms that are followed by fewer than ten analysts, recommendations are unanimous, meaning that all analysts submit identical firm ratings. This stands in stark contrast to firms that are followed by ten analysts or more, where consensus is achieved in less than 1% of cases. Such variation in consensus with analyst coverage is interesting in its own right and may warrant further investigation. In

Table 1

Number of analysts per rm in January 15, 2015 I/B/E/S vintage

	Under 10	Between 10 and 20	Between 20 and 30	Over 30
% of firms	71.18%	20.24%	7.01%	1.57%

Source: compiled by the author.

Table 2

Summary statistics for I/B/E/S consensus recommendations and changes in volatility of daily returns

	Mean	St. Dev.	Skew.	Kurt.	Max.	Min.
$^{R}i; T$	2.353	0.390	0.343	2.951	3.820	1.150
$V_i(21)$	0.123	0.428	2.291	14.312	4.122	–0.727
$V_i(63)$	–0.028	0.269	0.970	4.690	1.381	–0.582
$V_i(126)$	–0.025	0.233	1.364	9.364	2.020	–0.555
$V_i(252)$	0.202	0.296	1.198	6.842	2.183	–0.790

Source: compiled by the author.

our case, since we are interested in predictive abilities of securities analysts as a group, to ensure that mean recommendations that we use in-deed represent a meaningful consensus we restrict our attention to firms that are followed by ten analysts or more, of which there are 1,161 in the January 2015 I/B/E/S vintage.

3.2. Security Returns & Volatility

For all securities in our sample we use the adjusted daily closing prices from the Center for Research in Security Prices (CRSP) to estimate volatility changes one, three, six and twelve months following the date of recommendation issue. Since we only observe price data on trading days, this amounts to calculating $V_i(21)$, $V_i(63)$, $V_i(126)$, and $V_i(252)$, respectively which capture changes in volatility of underlying holding period returns that include returns arising from dividends and other distributions. *Table 2* shows summary statistics for our recommendations and volatility data. Histograms showing distributions of consensus recommendations and associated one-month volatility changes are provided in *Figure*. Next, we begin by conducting exploratory non-parametric analysis and proceed with the estimation of a formal copula model in Section 4.2.

4. RESULTS

4.1. Nonparametric Analysis: Empirical Copula Table

To gain an initial understanding of the nature of relationship between consensus recommendations and changes in return volatility for securities in our sample we begin by constructing a so-called empirical copula table for $R_{i,T}$ and $V_i(K)$, for different event horizons K. The *table* can highlight parts of the joint distribution range where association between the variables is strongest, and is often used as the initial step in the copula model selection process. For examples of applications of empirical copula tables see [19–21], among others. The construction of the table proceeds as follows. First, we obtain non-parametric estimates of the marginal models for recommendations and volatility as

$$\widehat{F}(x)=\frac{1}{n}\sum_{i=1}^{n}I\left(R_{i,T^*}\le x\right)\text{ and }\widehat{H}(y)==\frac{1}{n}\sum_{i=1}^{n}I\left(V_i\left(K\right)\le y\right), \quad (9)$$

where n = 1.161 is our sample size. For every observation j = 1,.., n, we next estimate corresponding normalized ranks as $\widehat{u_j}=\widehat{F}\left(R_{j,T^*}\right)$ and $\widehat{v_j}=\widehat{H}\left(V_j\left(K\right)\right)$.

Note that by construction u_j and v_j are uniformly distributed on [0;1] and additionally, when $R_{i,T}$ and $V_i(K)$ are stochastically independent, that is, when there is no relationship between consensus recommendations and changes in realized volatility, (u_j, v_j) are also jointly uniformly distributed on $[0;1]^2$. Deviations

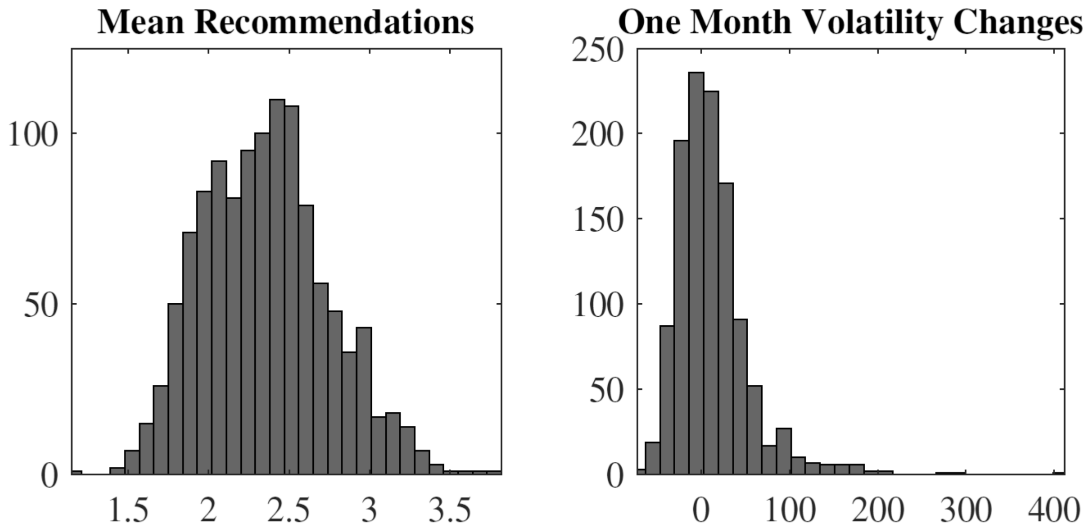

Fig. **Histograms for I/B/E/S consensus recommendations and associated one month volatility changes (V_i(21)). Distributions of V_i(63), V_i(126), and V_i(252) are of similar shape to that of V_i(21)**

Source: compiled by the author.

from uniformity indicate dependence, and the aim of the empirical copula table is to identify parts of the range such as, for example, distribution tails or certain distribution quad-rants, where clusters of (u_j; v_j) are present and therefore dependence exists. To construct the table, we sort estimated u_j and v_j in ascending order and allocate pairs (u_j; v_j) into 16 bins in accordance with rank. Letting $b_{k,\,s}$, k = 1,..,4, s = 1,..,4 denote bin count, the sorting is done so that each bin contains pairs (u_j, v_j) that lie between recommendations quartiles k and k — 1, and volatility change quartiles s and s — 1. That is, the first bin will contain $b_{1,1}$ pairs of observations that belong in the bottom 25% of recommendations and bottom 25% of volatility, and the last bin will contain $b_{4,4}$ observations that belong to the top 25% of recommendations and volatility changes respectively, and so on. When volatility changes are independent from recommendations, we can expect to see the same number of observations in every bin, which in our case amounts to approximately 73 observations per bin. Deviations from this count will indicate the tendency of $R_{i,\,T}$ and $V_i(K)$ to "cluster" together at a particular part of the joint distribution.

We collect empirical copula tables for $R_{i,\,T}$ and V_i(21), V_i(63), V_i(126), and V_i(252) in *Table 3*. For every event horizon, bold indicates the two bins with highest deviation from the count expected under stochastic independence of recommendations and volatility changes, which is 73 observations per bin. Interestingly, it appears that at all event horizons greater than one month, largest deviations from independence occur in bins (1,1) and (4,4), which are the upper-right and lower-left tails of the joint distribution. Such clusters suggest that the lowest 25% of consensus recommendations coincide with lowest 25% of volatility changes, while largest 25% of recommendations — with greatest 25% of volatility changes. In other words, it seems that best recommendations are associated with largest volatility declines, while worst recommendations — with largest jumps in realized volatility.

While these results are illustrative, they do not amount to a formal test for association between $R_{i,\,T}$ and $V_i(K)$. Next, using counts in *Table 3* as the starting point we select and t a copula model to our data and aim to formally test for the presence of statistically-meaningful clusters in distribution tails.

4.2. Semiparametric Analysis: Symmetrized Joe-Clayton Copula Model

Association that is skewed towards the tails of the distribution as we observe in *Table 3* is often studied using the so-called upper- and lower-tail dependence coefficients denoted λ_u and λ_l respectively and defined as

$$\lambda_u = \lim_{t \to 1^-} Pr\left[F(x) \geq t | G(y) \geq t \right] = \lim_{t \to 1^-} \frac{1 - 2t + C(t,t)}{1 - t}, \quad (10)$$

$$\lambda_l = \lim_{t \to 0^+} Pr\left[F(x) \leq t | G(y) \leq t \right] = \lim_{t \to 0^+} \frac{C(t,t)}{t}. \quad (11)$$

These parameters depend only on the copula C and show the limiting probability that one variable is larger or smaller than 100'th percentile, conditional on the other variable also exceeding or being below its 100'th

Table 3

Empirical copula changes for I/B/E/S consensus recommendations and one to twelve month changes in realized volatility. Bold highlights bins with the two largest deviations from the count expected under independence and volatility (73 observations)

	One month vol.					Three month vol.			
Bin	1	2	3	4	Bin	1	2	3	4
1	83	77	62	67	1	96	72	70	51
2	73	80	73	68	2	75	80	67	72
3	74	67	64	82	3	54	77	73	83
4	60	66	91	74	4	65	61	80	85
	Six month vol.					**Twelve month vol.**			
Bin	1	2	3	4	Bin	1	2	3	4
1	104	77	59	49	1	94	79	64	52
2	82	68	76	68	2	86	70	71	67
3	51	75	79	82	3	55	76	85	71
4	53	70	76	92	4	55	65	70	101

Source: compiled by the author.

percentile as t approaches 1 or 0. Larger values of either λ_u or λ_l indicate greater likelihood that large of small extremes of the variables will co-occur, which in our case amounts to association between best (smallest) recommendations and largest volatility declines, and worst (largest) recommendations and biggest jumps in volatility. For an excellent introduction to tail dependence see Section 5.4 of [15].

One family of copulas that is particularly well-suited for the analysis of dependence in both upper and lower distribution tails is the so-called Symmetrized Joe-Clayton Copula (SJC) of [14] defined as

$$C_{SJC}\left(u,v;k,r\right)=0.5C_{JC}\left(u,v;k,r\right)+ \\ +C_{JC}\left(1-u,1-v;k,r\right)+u+v-1, \qquad (12)$$

where $C_{JC}(u; v; k; r)$ is the Joe-Clayton (or BB 7) copula given by

$$C_{JC}\left(u,v;k,r\right)=1-\begin{pmatrix}1-\{\left[1-\left(1-u\right)^{k}\right]^{-r}+\\+\left[1-\left(1-v\right)^{k}\right]^{-r}-1\}^{-1/r}\end{pmatrix}^{1/k}, \qquad (13)$$

and k, $r \geqslant 0$ are parameters such that $k=1/\log_2\left(2-\lambda_u\right), r=-1\log_2\left(\lambda_l\right)$, $\lambda_u,\lambda_l\in\left(0,1\right)$. Estimates of tail dependence coefficients can therefore be obtained relatively easily by estimating 15 using IFM or CML. Since bounds on recommendations can complicate the specification of a parametric marginal model for *F* we focus on CML estimation here. We use bootstrap to obtain associated standard errors.

4.3. Estimation Results

Table 4 provides CML estimates of upper- and lower-tail dependence coefficients associated with the SJC copula in (15) obtained using analyst recommendations and corresponding changes in return volatility at one, three, six, and twelve-month horizons, or $K = 21$, $K = 63$, $K = 126$ and $K = 252$, respectively. We first estimate the model using the entire sample, where some ratings represent reiterations, some — upgrades, and some — downgrades of analysts' previous opinions. The resulting pooled estimates λ_u and λ_l measure unconditional upper- and lower-tail dependence between recommendations and volatility, which we report first. We then re-estimate the model separately for the securities which have been upgraded, downgraded, or revised in either direction relative to previous I/B/E/S vintage. Estimates λ_u and λ_l in these three cases therefore capture association between recommendations and volatility conditional on recommendation revisions.

Table 4

CML estimates of Symmetrized Joe-Clayton copula parameters for analyst recommendations and changes in return volatility. Parentheses contain t-ratios

	$\hat{\lambda}_u$	$\hat{\lambda}_l$
Unconditional		
One month volatility changes	0.0015 (0.7925)	0.0001 (0.0001)
Three month volatility changes	0.0019 (0.2186)	0.0203 (1.0584)
Six month volatility changes	0.0142 (0.6958)	0.0401 (1.5603)
Twelve month volatility changes	0.0085 (0.4851)	0.0399 (1.7608)
Following Upgrades		
One month volatility changes	0.0026 (0.0498)	0.0001 (0.0003)
Three month volatility changes	0.0336 (0.5881)	0.1510 (2.1541)
Six month volatility changes	0.0134 (0.2383)	0.1530 (1.8957)
Twelve month volatility changes	0.0001 (0.0001)	0.0987 (1.9153)
Following Downgrades		
One month volatility changes	0.0116 (0.3757)	0.0221 (0.6371)
Three month volatility changes	0.0325 (0.8680)	0.0313 (0.7287)
Six month volatility changes	0.0878 (1.4178)	0.0638 (1.2443)
Twelve month volatility changes	0.0024 (0.0840)	0.1177 (2.0942)
Following Any Revision		
One month volatility changes	0.0069 (0.3780)	0.0009 (0.0594)
Three month volatility changes	0.0314 (0.9487)	0.0062 (1.6155)
Six month volatility changes	0.0658 (1.3556)	0.0828 (1.8391)
Twelve month volatility changes	0.0001 (0.0008)	0.1136 (2.1250)

Source: compiled by the author.

We highlight estimates that are statistically-positive at 5% s.l. in *Table 4*. Interestingly, at all event horizons and for all sample groupings, estimate of the upper-tail dependence coefficient λ_u is not statistically-different from zero, suggesting that the clustering between high (that is, poor) recommendations and high subsequent volatility of security returns is not significant. On the other hand, estimate of the lower-tail dependence coefficient λ_l is significant in all sample groupings at twelve month event horizon, and at six month horizon when conditioning on recommendation upgrades and revisions, but not downgrades. It therefore appears that lowest (or best) recommendations indeed tend to be associated with smallest volatility changes, or largest volatility declines. In other words, it seems that best-rated stocks experience largest drops in volatility, but only at longer event horizons that are six to twelve months from recommendation issue. The converse does not seem to hold for worst-rated stocks.

It is further worth noting that such lower-tail dependence seems stronger in the conditional than unconditional case, indicating that association be-tween recommendations and volatility that we document in *Table 4* is more pronounced among stocks with revised

Table 5

Realized 5% Value-at-Risk for best-rated and worst-rated portfolios

	Best-Rated	Worst-Rated	Difference
Unconditional			
One month horizon	−3.06%	−2.90%	5.52%
Three months horizon	−1.97%	−2.43%	−18.93%
Six months horizon	−1.75%	−2.05%	−14.63%
Twelve months horizon	−2.33%	−2.39%	−2.51%
Following revisions			
One month horizon	−2.96%	−3.34%	−11.38%
Three months horizon	−1.97%	−2.91%	−32.30%
Six months horizon	−1.94%	−2.55%	−23.92%
Twelve months horizon	−2.29%	−2.98%	−23.15%

Source: compiled by the author.

Table 6

Changes in 1% VaR for best-rated and worst-rated portfolios after recommendation issue

	Best-Rated	Worst-Rated
Unconditional		
One month horizon	−3.43%	7.39%
Three months horizon	−24.90%	−10.34%
Six months horizon	−22.94%	−5.43%
Twelve months horizon	−11.84%	14.10%
Mean VaR Change	−15.78%	1.43%
Following revisions		
One month horizon	−13.79%	0.42%
Three months horizon	−26.67%	−28.83%
Six months horizon	−18.94%	−6.04%
Twelve months horizon	16.91%	16.32%
Mean VaR Change	−10.62%	−4.53%

Source: compiled by the author.

ratings. This is consistent with [7] and [8], among others, who also find revisions to be informative.

4.4. Application to Portfolio Risk Management

Next, we use our findings to assess the investment value of analyst recommendations in the applied context, focusing on the management of portfolio market risk. Estimates from the preceding section suggest that strongest linkages between recommendations and return volatility occur at the extremes of the recommendations spectrum. We use this observation to create two equally-weighted portfolios on the basis of I/B/E/S recommendations: a so-called "best-rated" portfolio containing top five percent of stocks in the January 2015 vintage, and a "worst-rated" portfolio containing bottom five percent of stocks by mean recommendation. We then estimate the realized daily 5% Value-at-Risk (VaR) for these portfolios during one-, three-, six-, and twelve-months periods after recommendation issue. We do not re-balance these portfolios during these intervals. As with copula parameter estimates, we first conduct this analysis for all securities in the sample and then repeat it separately only for the securities where recommendations have changed. We collect our estimates in *Table 5*.

Realized value-at-risk appears to be substantially lower in nearly-all cases for top-rated securities, except in the pooled sample at one month horizon. Also, this difference seems greatest when conditioning on recommendation changes, which further supports informativeness of recommendation revisions, now from the standpoint of portfolio risk.

Interpretation of these differences, however, is complicated by the possibility of inclusion of risk into the analysts' rating process. Systemic, long-term differences in levels of security risk that arise, for example, due to different levels of industry or operating risks may drive ratings in the first place and explain strati cation that we document in Table 5. In other words, it may be unsurprising to see lower risk levels be associated with best-rated stocks as analysts reserve most favorable ratings for stable, well-performing companies operating in safer markets and practicing good corporate governance, which leads to lower volatility of the share price in the long term.

To see whether analyst opinions contain predictive information that is distinct from the observable trend in company risk, we repeat our analysis involving best-rated and worst-rated portfolios, but now focus on changes in portfolio value-at-risk following recommendation issue. That is, we now seek to identify whether analysts can reliably predict changes in these risk trends. Such standardization amounts to conditioning of our estimates on past volatility of security returns and ensures that any differences in VaR between the two portfolios that we document are not driven by idiosyncrasies already embedded into recommendations. We collect our findings involving VaR changes in *Table 6*.

Interestingly, here once more value-at-risk appears to drop substantially for the portfolio consisting of best-rated securities at nearly all horizons following recommendation issue. This is in contrast to worst-rated stocks, where the picture is mixed and on average these stocks experience an increase in value-at-risk in the unconditional case, and a much-smaller mean drop in VaR across our selected time frames when conditioning on recommendation changes. Overall, our estimates remain in line with the differences documented in *Table 5* and now suggest that buying best-rated stocks may lead to substantial reduction in portfolio-level risk measures over time, which supports the notion that analysts' recommendations contain predictive information and hence positive investment value.

Lastly, it is worth noting that while it is perhaps not surprising to see that the reduction in volatility among best-rated securities that we document in Section 4.3 translates into a drop in portfolio value-at-risk, results reported in *Tables 5* and *6* quantify this effect using a widely-used portfolio-level risk metric and suggest that incorporating I/B/E/S ratings into the risk-management process may have a meaningful effect.

5. DISCUSSION

Empirical results collected in Section 4.2 appear to provide first evidence of linkages between recommendations issued by security analysts and subsequent realized volatility of stock returns in the literature. These links could exist for a number of reasons and establishing causality in that respect is beyond the scope of this work. A plausible explanation, however, could relate to analysts' ability to identify fundamental factors in company financials that are related to future price volatility that are reflected in recommendations. This presents scope for future work, where the marginals can now be made conditional on stock fundamentals.

Similarly, our findings in Section 4.4 seem to provide first and positive assessment of investment value of these recommendations from the standpoint of portfolio risk. Several related important questions remain unanswered, however, and represent scope for future work. First, since our focus here is on a cross-section of ratings collected in a single I/B/E/S vintage, it would be interesting to see how the ratings-risk relationship evolves over time, and in particular, examine its nature shortly before and during periods of financial crisis. Second, analysts' predictive ability with respect to future levels of risk may be concentrated in certain sectors, and its distribution among various industries may be of interest. It could also be interesting to see whether any analyst characteristics such as, for example, years of experience or reputation and size of the rm are predictive of their ability to forecast risk. Lastly, from the standpoint of market efficiency it may be interesting to examine derivative trading strategies which aim to utilize the predictive power of analyst recommendations identified here.

REFERENCES

1. Stickel S. The anatomy of the performance of buy and sell recommendations. *Financial Analysts Journal*. 1995;51(5):25–39. DOI: 10.2469/faj.v51.n5.1933
2. Womack K.L. Do brokerage analysts' recommendations have investment value? *The Journal of Finance*. 1996;51(1):137–167. DOI: 10.1111/j.1540–6261.1996.tb05205.x
3. Barber B., Lehavy R., McNichols M., Trueman B. Can investors profit from the prophets? Security analyst recommendations and stock returns. *The Journal of Finance*. 2001;56(2):531–563. DOI: 10.1111/0022–1082.00336

4. Jegadeesh N., Kim W. Value of analyst recommendations: International evidence. *Journal of Financial Markets*. 2006;9(3):274–309. DOI: 10.1016/j.finmar.2006.05.001
5. Green T.C. The value of client access to analyst recommendations. *The Journal of Financial and Quantitative Analysis*. 2006;41(1):1–24. DOI: 10.1017/S 0022109000002404
6. Ivković Z., Jegadeesh N. The timing and value of forecast and recommendation revisions. *Journal of Financial Economics*. 2004;73(3):433–463. DOI: 10.1016/j.jfineco.2004.03.002
7. Barber B.M., Lehavy R., Trueman B. Ratings changes, ratings levels, and the predictive value of analysts' recommendations. *Financial Management*. 2010;39(2):533–553. DOI: 10.1111/j.1755–053X.2010.01083.x
8. Jegadeesh N., Kim J., Krische S.D., Lee C.M.C. Analyzing the analysts: When do recommendations add value? *The Journal of Finance*. 2004;59(3):1083–1124. DOI: 10.1111/j.1540–6261.2004.00657.x
9. Medovikov I. Can analysts predict rallies better than crashes? *Finance Research Letters*. 2014;11(4):319–325. DOI: 10.1016/j.frl.2014.08.001
10. Loh R.K., Stulz R.M., 2011. When are analyst recommendation changes influential? *The Review of Financial Studies*. 2011;24(2):593–627. DOI: 10.1093/rfs/hhq094
11. Devos E., Hao W., Prevost A.K., Wongchoti U. Stock return synchronicity and the market response to analyst recommendation revisions. *Journal of Banking & Finance*. 2015;58:376–389. DOI: 10.1016/j.jbankfin.2015.04.021
12. Sklar A. Fonctions de répartition à n dimensions et leurs marges. Publications de l'Institut de Statistique de l'Université de Paris. 1959;8:229–231.
13. Patton A.J. Modelling asymmetric exchange rate dependence. *International Economic Review*. 2006;47(2):527–556. DOI: 10.1111/j.1468–2354.2006.00387.x
14. Joe H. Multivariate models and multivariate dependence concepts. Boca Raton, FL: CRC Press; 1997. 424 p. (Monographs on Statistics and Applied Probability. Vol. 73).
15. Nelsen R.B. An introduction to copulas. New York: Springer-Verlag; 2006. 272 p. (Springer Series in Statistics).
16. Cherubini U., Luciano E., Vecchiato W. Copula methods in finance. Chichester: John Wiley & Sons Ltd.; 2004. 310 p. (Wiley Finance Series).
17. Patton A.J. Copula-based models for financial time series. In: Mikosch T., Kreiß J.P., Davis R., Andersen T., eds. Handbook of financial time series. Berlin, Heidelberg: Springer-Verlag; 2009:767–785.
18. Joe H., Xu J.J. The estimation method of inference functions for margins for multivariate models. The University of British Columbia. Technical Report. 1996;(166). URL: https://open.library.ubc.ca/cIRcle/collections/facultyresearchandpublications/52383/items/1.0225985 (accessed on 11.01.2019).
19. Knight J., Lizieri C., Satchell S. Diversification when it hurts? the joint distributions of real estate and equity markets 1. *Journal of Property Research*. 2005;22(4):309–323. DOI: 10.1080/09599910600558520
20. Ning C. Dependence structure between the equity market and the foreign exchange market — A copula approach. *Journal of International Money and Finance*. 2010;29(5):743–759. DOI: 10.1016/j.jimon n.2009.12.002
21. Medovikov I. When does the stock market listen to economic news? New evidence from copulas and news wires. *Journal of Banking & Finance*. 2016;65:27–40. DOI: 10.1016/j.jbank n.2016.01.004

ABOUT THE AUTHOR

Ivan S. Medovikov — Associate Professor at the Department of Economics, Brock University, Ontario, Canada. Dr. Medovikov's areas of concentration include financial econometrics and empirical finance. He held a previous research role at the Bank of Canada and currently serves as a research director at Spartan Fund Management, Inc., a Toronto-based alternative asset management firm, and is a managing partner of Price Street, Inc., a financial research consultancy
imedovikov@brocku.ca

ИНФОРМАЦИЯ ОБ АВТОРЕ

Иван Сергеевич Медовиков — доцент кафедры экономики Университета Брок, Онтарио, Канада. Сферы научных интересов включают финансовую эконометрику и эмпирические финансы. В прошлом работал научным исследователем в центральном банке Канады. На данный момент также является директором по исследовательской деятельности в Spartan Fund Management, Ind.— инвестиционной компании в Торонто. Управляющий партнер в Price Street, Inc., занимающейся финансовым консалтингом
imedovikov@brocku.ca