Co-movement of revisions in short- and long-term inflation expectations

António Armando Antunes
Banco de Portugal and NOVA SBE

May 2015

Abstract
This article studies the co-movement between large daily revisions of short- and long-term inflation expectations using copulas. The main findings are: first, the co-movement between unusually large changes in short- and long-term inflation expectations increased markedly since mid-2012, which implies that long-term inflation expectations might not be, in a precise sense, well-anchored. Second, this co-movement measure is quite noisy. Finally, the result is shown not to be an artifact of the methodology or of the specific data used in the analysis. (JEL: C14, C46, G12)

Introduction

Market-based inflation expectations are widely used by market participants and policymakers for decision making and for inferring the likely monetary policy decisions of central banks. The alternative survey-based inflation expectations are also widely used but, for the purposes of this article, are not suitable given the lower frequency of available data. Market-based inflation expectations can be determined in several ways but perhaps the most popular method resorts to market prices of zero-coupon inflation swaps. These financial instruments are composed of a fixed leg and a variable leg and can be used to wedge against inflation fluctuations. For example, suppose that market participant A wants to insure herself against inflation fluctuations for holding a nominal asset for a period of five years starting from now. She can enter a zero-coupon inflation swap contract in the following terms: at the end of the five years, she receives the actual change in the relevant inflation index, which in the euro area can be for example the HIPC excluding tobacco, times the notional amount of the contract. This exactly compensates her for the changes of opposite sign in the real value of the nominal asset. At the same time, she pays the fixed leg of the contract to counter-party B, which is determined using the fixed rate

Acknowledgements: I thank Ildeberta Abreu, Rafael Barbosa, Nikolay Iskrev and Paulo Rodrigues, as well as participants in internal seminars, for useful discussions and help.

E-mail: aantunes@bportugal.pt

compounded for five years. Only one cash flow is exchanged at maturity, but the position can be closed at any moment by selling the contract in the market. The rate of the fixed leg of the contract is the expected inflation rate for the next five years. In fact, B enters the contract only if she believes that the fixed leg rate is going to be at least the actual inflation rate at maturity. On the other hand, A enters the contract only if she believes that the actual inflation is going to be at least the fixed leg rate. Of course there are additional effects involved. In particular, because A is effectively wedged against inflation risk, B has to be compensated through an inflation risk premium.

Using market-based inflation rate expectations, this article assesses the co-movement between daily revisions in short- and long-term inflation expectations using copulas, a special class of multivariate distribution functions. The main advantage of using copulas lies in their simple implications in terms of dependence of random variables, especially in the tails of the distribution. This allows for an assessment of the degree with which changes in long-term inflation expectations co-move with large swings in short-term inflation expectations. Moreover, certain copulas allow one to distinguish between upward and downward revisions in expectations.

Policymakers often mention that long-term inflation expectations are "well-anchored". However, this expression can mean different things. Sometimes it refers to the fact that the level of expectations is hovering close to a commonly accepted target level. On other occasions, the expression simply asserts that revisions of short-term inflation expectations should not per se imply revisions of long-term inflation expectations. One implication of this is that revisions in short- and long-term inflation expectations should not co-move significantly. The two meanings are not equivalent and have distinct implications in terms of the suitable methods for investigating whether inflation expectations are well-anchored. While the first focuses on levels and calls for a more traditional time series analysis, the second suggests using methods with an emphasis on correlation and co-movement, and not necessarily keeping track of the level of the inflation expectations. This article adopts the second type of approach. Moreover, special attention is paid to large innovations in inflation expectations, as these are more likely to represent fundamental changes in expectations than normal market fluctuations of smaller magnitude.

In a world where the central bank is deemed credible by market participants and with perfectly anchored long-term inflation expectations, one would expect that large revisions of short-term inflation expectations displayed little co-movement with large revisions of long-term inflation expectations. For example, a sudden oil price drop implying a large revision downwards of the short-term inflation expectations should not imply a revision of the same magnitude (in relative terms) in long-term inflation expectations.

Likewise, if one observes large revisions in long-term inflation expectations when there are large revisions in the short-term expectations, then the idea that long-term inflation expectations are solidly anchored becomes less obvious. In the limit, if one were to observe a one-to-one co-movement between these two measures, surely inflation expectations would not be anchored: they would be reacting immediately and significantly to the same information that produced swings in short-term expectations, with potentially highly disruptive effects in the effectiveness of monetary policy.

There is a relatively large literature on this topic which uses high frequency data and focuses mostly on the effects of news on long-term inflation expectations. This literature usually looks at the possibility of occurrence of structural breaks in a context of regression analysis (see, for example, Gürkaynak *et al.* 2010; Galati *et al.* 2011; Nautz and Strohsal 2015). In this article it is assumed that news are incorporated both in short- and long-term expectations but, if long-term inflation expectations are well-anchored, the effect on them would be small, whereas the effect on the short-term ones would be large. This should induce a low degree of co-movement between inflation expectations at long and short horizons. Using estimated copulas, it is shown that co-movement between changes in short- and long-term inflation expectations increased since 2012. This is in contrast with the absence of any significant co-movement in the previous low inflation period of end-2009. Moreover, these effects are shown not to be an artifact of the data, as simulations with random permutations of data eliminate them. Tail dependence between revisions in short- and lagged long-term inflation expectations persists but only for lags of one or two days, especially in the most recent portion of the sample. Finally, different choices for short- and long-term inflation expectations do not change the results in any meaningful way. While noisy, the observed co-movement in large swings suggests an increasing likelihood that long-term inflation expectations might have become de-anchored.

Inflation expectations and co-movement

In this article inflation expectations are taken from zero-coupon inflation swap rates. In terms of notation, average inflation prevailing from now until five years from now, for example, is denoted by π_{5y0y}, average inflation prevailing from next year for the following three years is π_{3y1y}, and average inflation prevailing five years from now for the following five years is π_{5y5y}. There are restrictions among these values, and these restrictions allow one to compute all relevant expectations based only on zero-coupon inflation swap rates. For instance, if market participants are risk neutral in perfectly competitive and frictionless markets, the equality $(1+\pi_{5y0y})^5 = (1+\pi_{2y0y})^2(1+\pi_{3y2y})^3$ must hold. Notice how the two zero-coupon rates can be used to estimate π_{3y2y}.

Another example: $(1+\pi_{5y5y})^5 = (1+\pi_{4y5y})^4(1+\pi_{1y9y})$ must hold. In this article, the value of the short-term inflation expectation will be the expected inflation one year ahead for one year (π_{1y1y}), and the long-term inflation expectation measure will be defined in the period five years ahead for five years (π_{5y5y}).

Data

Data are daily from Bloomberg and span the period from 22Jun04 until 17Feb15. Figure 1 presents the evolution over time of the two chosen variables, π_{1y1y} and π_{5y5y}, as well as observed inflation at monthly frequency. Table 1 presents summary statistics of the levels and first differences of π_{1y1y} and π_{5y5y}, along other variables (see below). The first differences correspond to the daily revisions of long- and short-term inflation expectations and constitute the focus of this article. Table 2 displays contemporaneous correlations among these variables.

Variable	Obs.	Mean	Std. Dev.	Min.	Max.	Autocorr.
π_{1y1y}	2781	1.787	0.504	0.293	3.751	0.978
π_{5y5y}	2781	2.304	0.205	1.483	2.803	0.987
$\triangle\pi_{1y1y}$	2780	-0.001	0.105	-1.163	1.132	-0.418
$\triangle\pi_{5y5y}$	2780	0.000	0.033	-0.196	0.220	-0.267
x	2780	-0.005	0.999	-5.968	11.334	-0.065
y	2780	0.000	1.000	-7.368	5.507	-0.019
u	2780	0.500	0.289	0.000	1.000	-0.005
v	2780	0.500	0.289	0.000	1.000	0.028

TABLE 1. Summary statistics. Daily data for period 22Jun04–17Feb14. π_{1y1y} and π_{5y5y} are market-based inflation rates one year from now during one year and five years from now during 5 years, respectively, and $\triangle\pi_{1y1y}$ and $\triangle\pi_{5y5y}$ are first differences; x and y are $\triangle\pi_{1y1y}$ and $\triangle\pi_{5y5y}$ filtered through an AR(1) process for the conditional mean and a GARCH(1,1) for the conditional variance; u and v correspond to the empirical quantiles of variables x and y, respectively. Values for π_{1y1y}, π_{5y5y}, $\triangle\pi_{1y1y}$ and $\triangle\pi_{5y5y}$ in percentages, except autocorrelations. Values for x, y, u, v and autocorrelations in natural units.

Sources: Bloomberg and author's calculations.

From the summary statistics it is readily seen that, historically, short-term inflation expectations have lower mean and higher volatility than long-term inflation expectations. In first differences, this behavior carries through for volatility but not for the mean, as expected. Level variables have strong persistence. In first differences there is negative autocorrelation, suggesting that increases in inflation expectations are often followed by corrections in the next trading day. Contemporaneous correlation between revisions of short- and long-term inflation expectations is only -0.007.

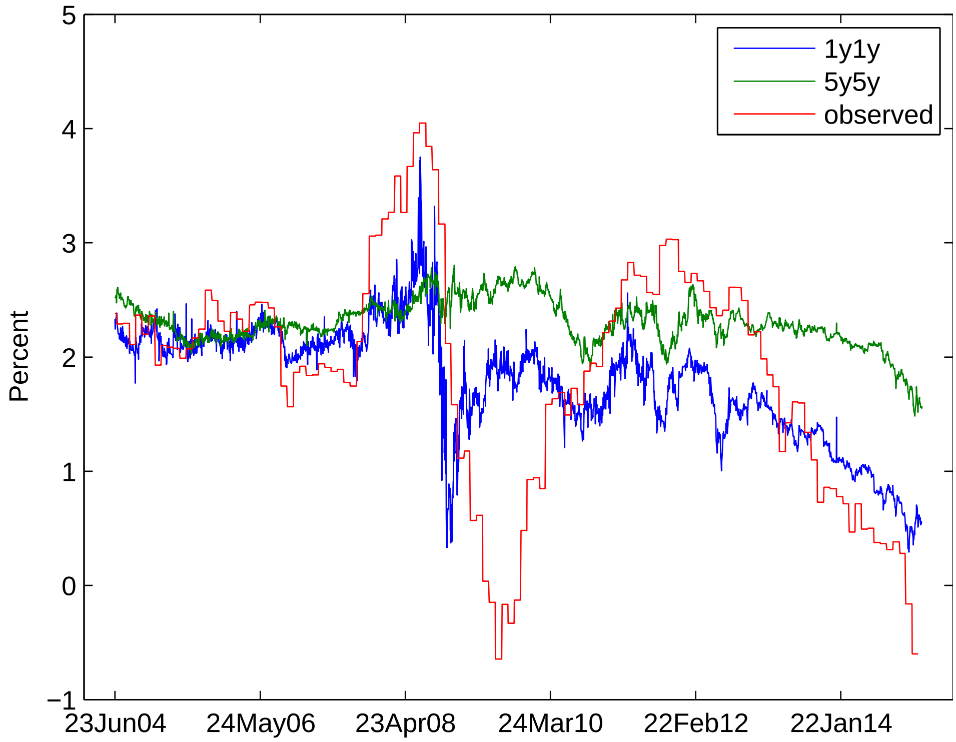

FIGURE 1: Market-based inflation rate expectations and observed inflation. Daily data for period 22Jun04–17Feb15. All values in percentage.

Source: Bloomberg.

	$\triangle\pi_{1y1y}$	$\triangle\pi_{5y5y}$	x	y	u	v
$\triangle\pi_{1y1y}$	1					
$\triangle\pi_{5y5y}$	-0.007	1				
x	0.761	0.047	1			
y	0.028	0.915	0.094	1		
u	0.681	0.049	0.893	0.088	1	
v	0.024	0.857	0.089	0.931	0.097	1

TABLE 2. Correlation matrix. Daily data for period 22Jun04–17Feb14. See legend of Table 1 for definitions of variables.

Sources: Bloomberg and author's calculations.

Conditional tail dependence

The study of co-movement between two random variables X and Y can be done in various ways. The first would be a simple correlation. This measure between -1 and 1 computes how X and Y co-move around their respective means. Sometimes this measure is enough for one's purposes. For example, the co-movement between two gaussian variables can be fully characterized by correlation. One problem with correlation as a measure of cross dependence is the fact that zero cross correlation does not in general imply independence. For example, the cross correlation between a random variable and its square is zero but they are clearly not independent. This in fact is a valid reason for not using correlation (or a linear regression coefficient) to study essentially unknown dependence among variables. Another problem with correlation is that it cannot be defined for certain distributions with heavy tails, as often is the case with financial returns (for examples of such distributions, see Resnick 2007).

An alternative way of studying co-movement between two variables is *conditional tail dependence*, and this is the focus of this article. To understand the notion of conditional tail dependence it is necessary first to define quantiles of a random variable. Quantile k of a random variable X is the value such that the probability of a random draw from X being less than or equal to that number is k. For example, the quantile 0.5 of a random variable is its median, and the interval defined by quantiles 0.025 and 0.975 is the 95% confidence interval for that random variable.

The idea of conditional tail dependence is simple: take values of variable X above a certain quantile k; compute the probability that the corresponding values of variable Y are above Y's quantile k; take the limit as k goes to 1. This is the so-called *upper tail dependence*. A similar measure can be computed for *lower tail dependence*, but in this case the limit is taken when k goes to 0. Intuitively, *this amounts to measuring the co-movement of two variables whenever either of them displays unusually large fluctuations*.[1]

This measure can be computed given the cumulative joint distribution function of the two variables, a function denoted by F. This function specifies the probability that a random realization of the two variables has both elements below the respective argument of F, so that for example $F(2,1)$ is the probability that, in a random draw from the joint distribution of X and Y, the draw from X is lower than 2 and the draw from Y is lower than 1. The marginal cumulative distribution is the cumulative distribution of one of the variables unconditional on the other; for example $F_X(x) = F(x, +\infty)$ would be the marginal cumulative distribution of X.

One way to proceed would be to estimate some parametric form for distribution F and then compute tail dependence. In practice, however, such an estimation is difficult and suffers from frequent scale and domain problems in terms of variables X and Y. An easier route to compute conditional tail dependence is using copulas.

1. See Appendix A for formal definition of tail dependence.

Copulas: intuition

Copulas are a special class of cumulative distribution functions; see Patton (2006b) for the etymology of this designation and a rationale for the use of copulas in practical applications, and Nelsen (2006) and Patton (2012) for a detailed exposition of the theory and practical aspects of copulas. The distinguishing features of a copula are two: first, its underlying random variables are defined in the $[0, 1]$ interval; second, its marginal distributions are those of an uniform distribution. Using a copula involves specifying marginal cumulative distribution functions of each random variable along with a function (that is, the copula) that connects them. In this way, the researcher can separate the modeling of the marginal distributions from the dependence between the two variables. The copula specification implies a certain shape for the dependence between the marginal distributions. In the case where the copula is the product of the two marginal cumulative distribution functions, the two variables are independent and one can separately estimate each marginal. Otherwise, one can efficiently resort to the estimation of the joint distribution using a copula. Since the copula captures dependence structures for any shape of the marginal cumulative distribution functions,[2] the copula approach to modeling related variables can be very useful from an estimation perspective.[3]

Data transformation

As with many distribution functions, copulas can be fitted to the data using maximum likelihood methods. However, inflation rate expectations do not necessarily have to lie on the interval between 0 and 1, as required by copulas, nor do they exhibit temporally uncorrelated behavior. In order to clean up data, in this analysis the original data, π_{1y1y} and π_{5y5y}, will be transformed in three steps. First, the variables of interest (daily revisions) are obtained by computing the first differences of the levels, yielding $\triangle\pi_{1y1y}$ and $\triangle\pi_{5y5y}$.

Second, because the sole interest of the analysis is dependence between variables, and to avoid spurious dependence stemming from persistence or heteroscedasticity, the resulting variables are filtered through an AR(1) model for the conditional mean and a GARCH(1,1) specification for the variance (for a similar approach, see for example Christoffersen *et al.* 2012). This yields standardized daily revisions in inflation expectations x and y, respectively for $\triangle\pi_{1y1y}$ and $\triangle\pi_{5y5y}$.

2. In fact, the dependence between the two distributions is, using a copula, invariant to monotonic transformations of the two random variables.

3. For a brief exposition of basic copula theory, as well as the notion of a dynamic copula, see Appendix B.

Third, standardized daily revisions in inflation expectations are mapped into numbers between 0 and 1 so that the resulting variables can be used to fit a copula. This is done through the computation of an empirical marginal cumulative distribution function. More specifically, take the time series of, say, the standardized revisions in inflation expectations one year ahead for one year, that is, the collection $\{x_t\}_{t=1,\ldots,T}$. Then there is a certain empirical marginal cumulative distribution function $\tilde{F}_X$ so that $u_t = \tilde{F}_X(x_t)$. (This function is an empirical, non-parametric counterpart to F_X.) Do a similar procedure for the long-term inflation expectations, y. Figure 2 represents the two empirical distribution functions. Variables u and v thus obtained are by construction approximately uniformly distributed.

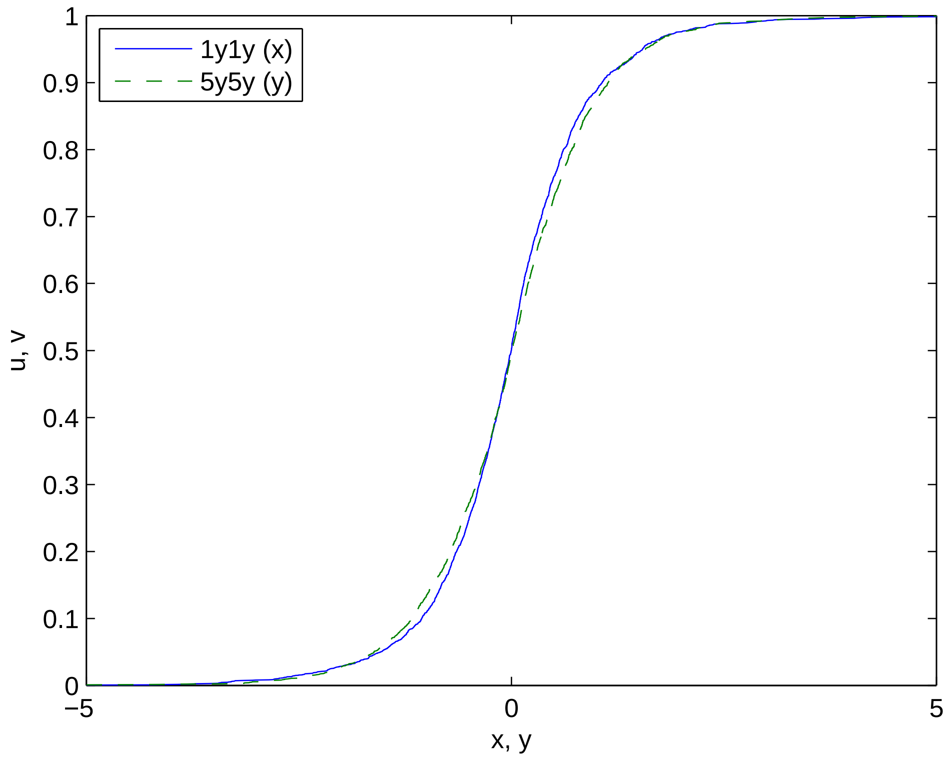

FIGURE 2: Empirical cumulative marginal distribution functions of x and y, the revisions of inflation expectations 1y1y and 5y5y standardized through applying an AR(1) conditional mean model and a GARCH(1,1) conditional variance model to the daily revisions of level variables.

Sources: Bloomberg and author's calculations.

Figures 3 and 4 present the daily innovations in inflation expectations, the standardized series and the uniform variables for the two variables of interest. Notice that there is substantial heteroscedasticity in both $\triangle\pi_{1y1y}$ and

$\triangle\pi_{5y5y}$, even though the latter exhibits less volatility, as previously seen. Heteroscedasticity is effectively removed by applying the filter mentioned above in both variables. Finally, the uniform transformations of x and y exhibit the expected behavior. Figure 5 shows a detail (observations during 2014) of the evolution of x and y.

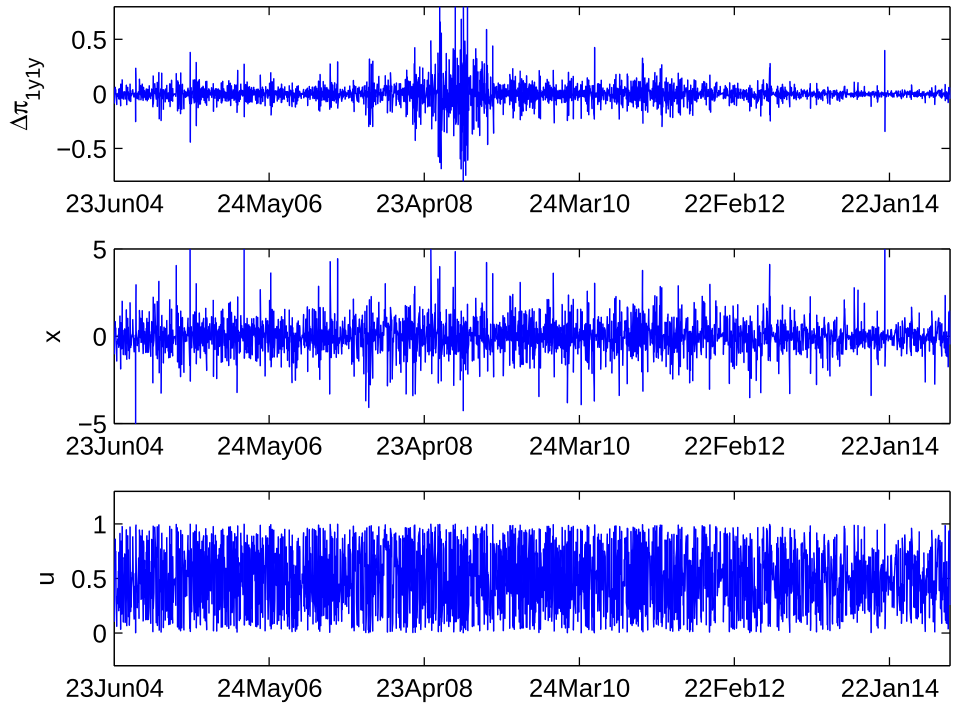

FIGURE 3: Evolution of $\triangle\pi_{1y1y}$, x and u. See legend of Table 1 for definitions of variables.

Sources: Bloomberg and author's calculations.

Going back to Tables 1 and 2, it can be seen that autocorrelation is mostly removed through the application of the AR(1) and GARCH(1,1) filters to the first differences of inflation expectations. Moreover, revisions of short- and long-term variables display relatively low contemporaneous correlation: the highest is u with v (0.097).

Results

The analysis consists of estimating several types of copulas in rolling windows of roughly one year, at the beginning of each quarter, and computing the

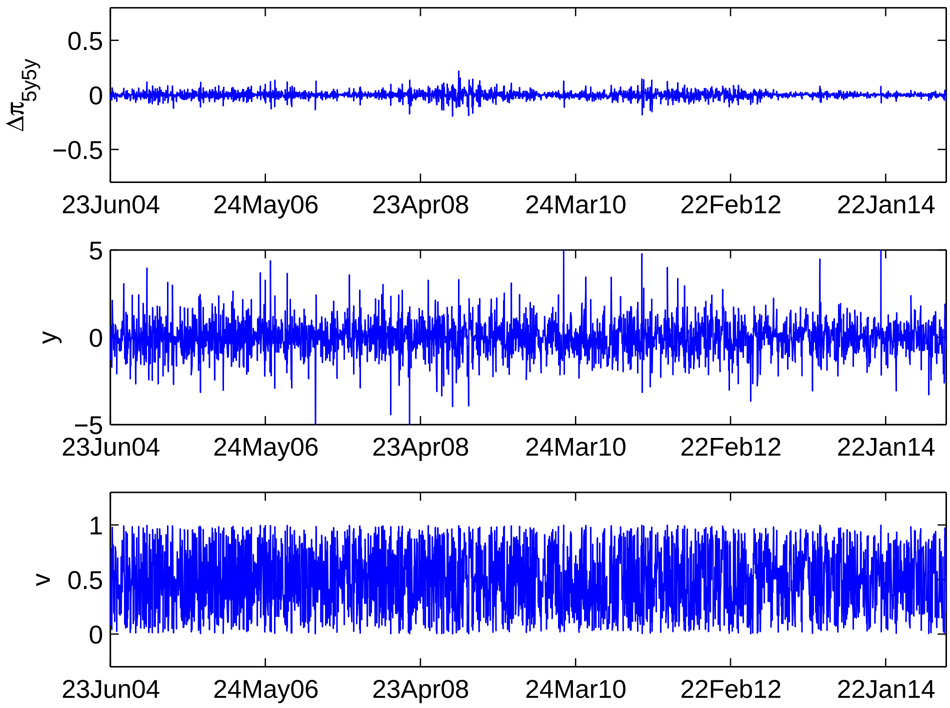

FIGURE 4: Evolution of $\triangle\pi_{5y5y}$, y and v. See legend of Table 1 for definitions of variables.

Sources: Bloomberg and author's calculations.

implied tail dependence. The estimated copulas differ in their parametric functional forms and, hence, in their characteristics in terms of symmetry and tail dependence.[4] A set of additional exercises and tests was also conducted but will only be briefly mentioned here.

Before looking at the evolution of tail dependence, a selection procedure was followed in which several different copulas were estimated. See Trivedi and Zimmer (2005) and Patton (2004, 2006a,b) and references therein for full descriptions of each copula. Table 3 summarizes the results. The ranking criterion was the number of times a copula is the best performer in each of the 39 quarters of the sample as measured by the value of its likelihood function. Under this criterion, the Student's t copula is the best performer, followed by

4. See Appendix B for a parametric example of a copula and references therein for full descriptions of copulas used in this section.

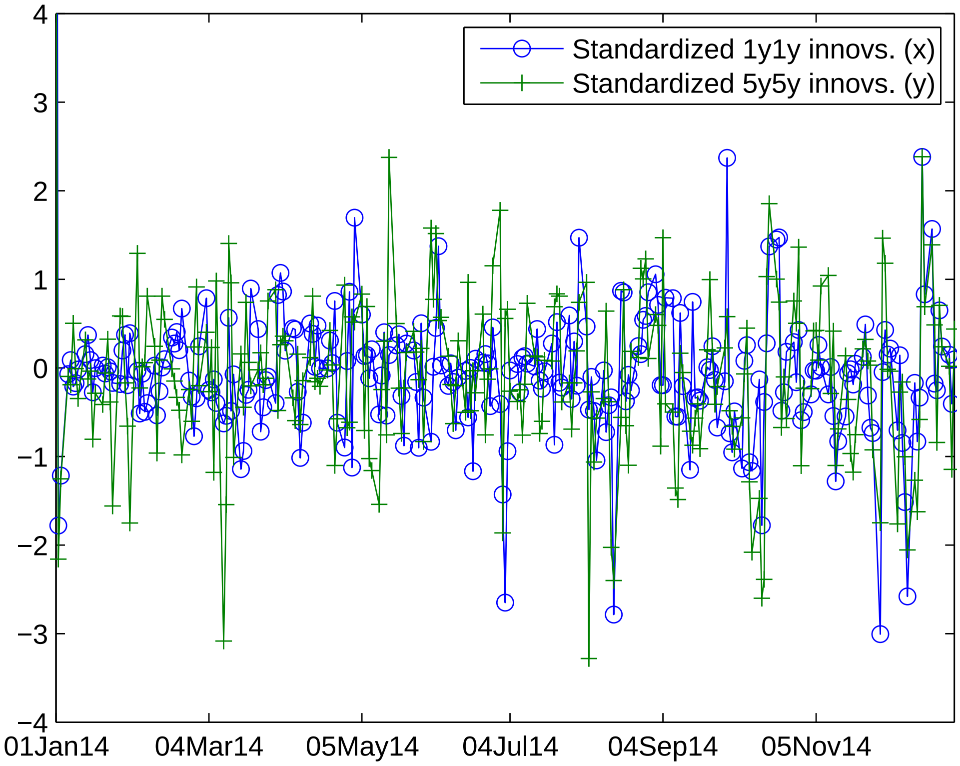

FIGURE 5: Evolution of x and y during 2014. See legend of Table 1 for definitions of variables.

Sources: Bloomberg and author's calculations.

the Normal, the Symmetrized Joe-Clayton (SJC), the Gumbel and the Rotated Gumbel.

At the beginning of each quarter, a copula was estimated using the available data of the previous 350 calendar days. The results are presented in Figures 6–8. The shaded areas are 90 percent confidence bands obtained through a bootstrap procedure (see Patton 2012). Looking at the results of Student's t copula (Figure 6), two features stand out. First, tail dependence is a noisy measure. The results are noisy and this volatility of the measure is still visible in the quarterly estimations reported in the figure.

The second salient aspect is that tail dependence increased markedly towards the end of the sample. The start of the increase in tail dependence can be traced back to 2012. The average tail dependence until 12Q3 was 0.011, and from 12Q4 on was 0.138. This is in stark contrast with the absence of

Copula	Tail dependence	# of quarters in which it was best
Student's	Yes, symmetric	20
Normal	No	9
Symmetrized Joe-Clayton (SJC)	Yes	7
Gumbel	Yes, upper tail	3
Rotated Gumbel	Yes, lower tail	0

TABLE 3. Ranking of estimated copulas according to the number of quarters that each copula performs the best.

Source: author's calculations.

any significant tail co-movement during the low inflation period of end-2009, when a fall in oil prices induced a marked decrease in inflation.

The figure also depicts the correlation parameter.[5] While at first tail dependence is fairly small, there is a period when, while there is correlation between the two series, the distribution becomes approximately Normal and no tail dependence occurs. After that, in 12Q4, tail dependence starts increasing consistently.

Among the copulas displaying tail dependence, the second best performer is the SJC copula and from the results depicted in Figure 7 one can see that upper tail dependence was higher than lower tail dependence during most of the sample. This means that large positive revisions in short- and long-term inflation expectations were more likely to be associated than large negative revisions. Towards the end of the sample (14Q2) lower tail dependence increases markedly. It should be noted that, since highly volatile data are being used, the distinction between upward revisions and downward revisions is not so clear-cut as with, say, quarterly data. Indeed, even when there seems to exist a secular trend to lower inflation, when one looks at longer spans of time (like, for example, during 2014 in Figure 1) daily filtered data still looks like white noise (see for example the filtered series in Figure 5), as expected, and there are as many upswings as there are downswings.

For the Gumbel copula, tail dependence decisively exceeds the 0.1 mark from 12Q3 on, and climbs to 0.4 towards the end of the sample. The Rotated Gumbel results are similar and hence not shown.

The Normal copula also performs well, although it has zero tail dependence. That is not surprising because the Student's t copula (which nests

5. Student's t copula estimation involves two parameters: correlation and degrees of freedom. When the estimated degrees of freedom of the copula become large, the copula converges to the Normal copula and there is no tail dependence.

the Normal copula as a particular case) has many degrees of freedom in many quarters, and this makes it very similar to the Normal copula in those quarters.

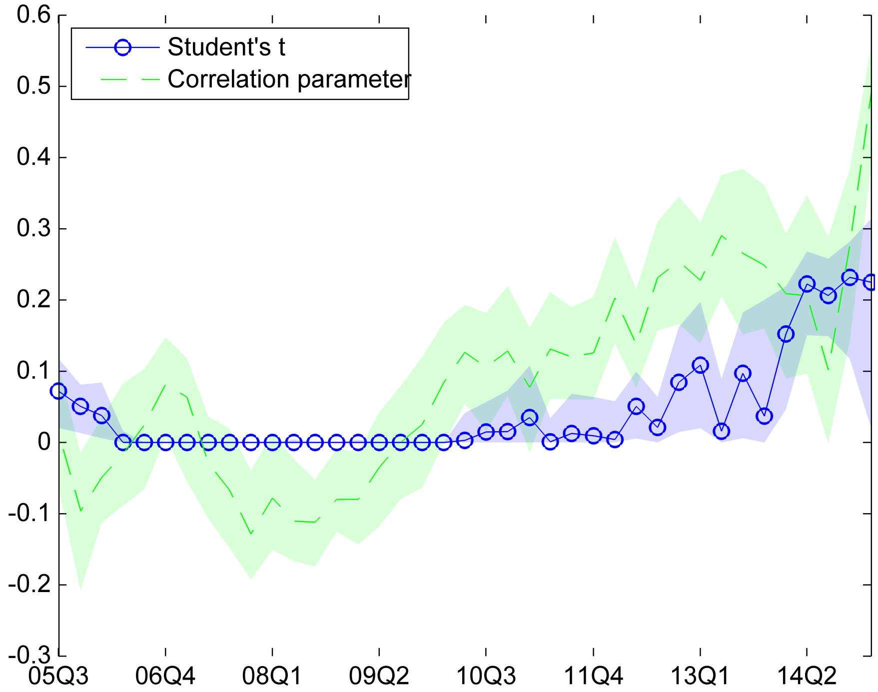

FIGURE 6: Tail dependence using estimated Student's t copulas at the beginning of each quarter using data from the previous 350 calendar days.

Source: author's calculations.

The general conclusion of the exercise is that the increase in tail dependence is very sharp since late 2012.

Additional exercises

Three additional exercises were performed.[6] The first is a robustness check where the whole procedure is repeated with a random permutation of time series $\{y_t\}_{t=1,\ldots,T}$ instead of the original series. The idea is to check whether there are artifacts of the data not related to co-movement that induce tail dependence. Given that the permutation should destroy all the time and cross dependence, one should observe essentially no tail dependence between the

6. Detailed results available upon request.

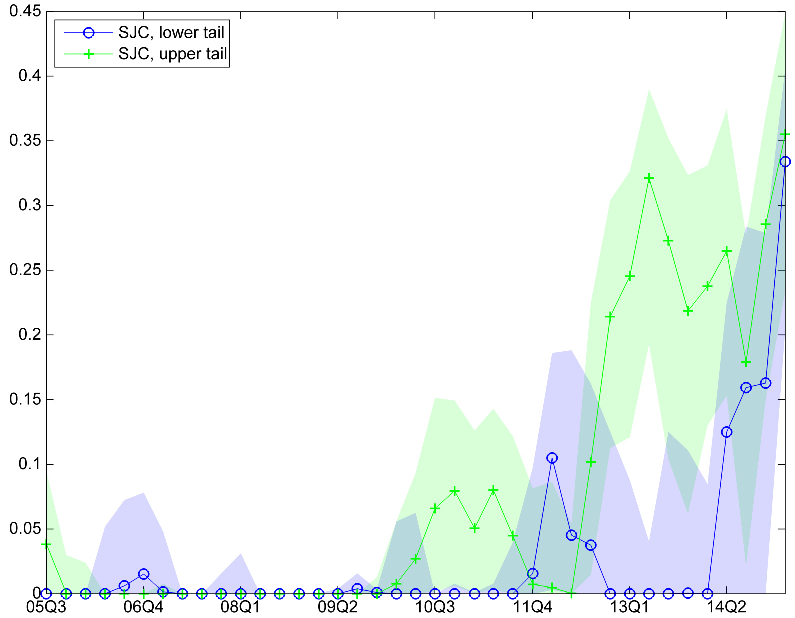

FIGURE 7: Tail dependence using estimated Symmetrized Joe-Clayton copulas at the beginning of each quarter using data from the previous 350 calendar days.

Source: author's calculations.

two variables. Indeed, the results show very low tail dependence throughout. The tail dependence parameters are found to be essentially zero. The second exercise was to perform the analysis with lags of one day and five days (which for this data set is one week) in variable y. The results for a 1-day lag display co-movement, although at a smaller level than the original estimates and concentrated in the final part of the sample. The co-movement dies out very fast and at a one-week lag it essentially has disappeared. All in all, this exercise suggests that there is time tail dependence at very short lags. The third exercise was to perform the analysis with different measures of short- and long-term inflation expectations, such as π_{2y1y} and π_{3y5y}. The results, however, remain essentially unaltered.

Concluding remarks

This article addresses the question of co-movement between revisions of short- and long-term inflation expectations. In particular, it focuses on a

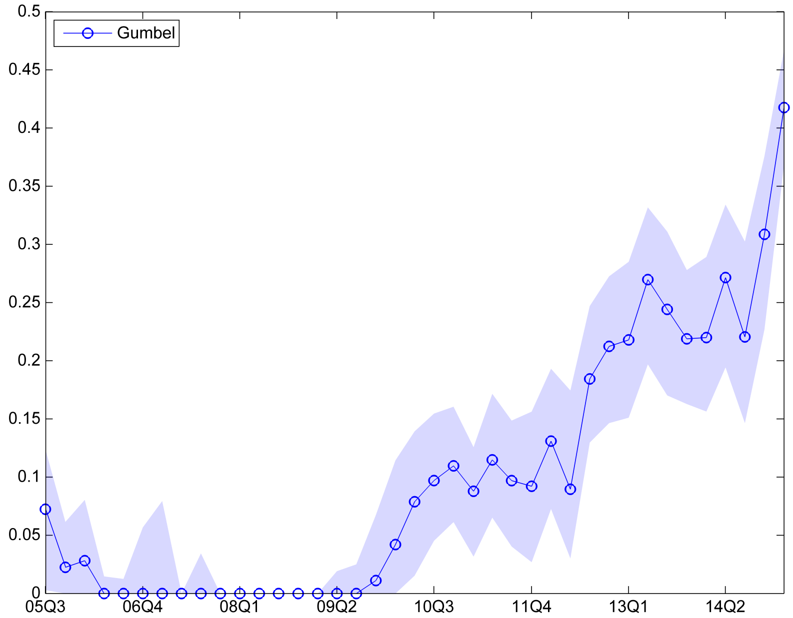

FIGURE 8: Upper tail dependence using estimated Gumbel copulas at the beginning of each quarter using data from the previous 350 calendar days.

Source: author's calculations.

measure called tail dependence, which looks at the probability that the two variables co-move when relatively large changes occur in one of them. Under the particular interpretation that inflation expectations are well-anchored when large innovations in short- and long-term inflation expectations do not co-move, this article shows that the case for well-anchored inflation expectations is not as strong since mid-2012 as it was before. This result is robust to different definitions of short- and long-term inflation expectations and does not seem to be an artifact of the data, produced for example by persistence or heteroscedasticity, and rapidly fades away when the data are not synchronous. Further work would include investigating the possibility of structural breaks in tail dependence in the context of copulas, and assessing the direction of causality, if any, in co-movement.

References

Braun, Valentin and Martin Grziska (2011). "Modeling Asymmetric Dependence of Financial Returns with Multivariate Dynamic Copulas." Mimeo, Goethe Universität, Frankfurt, and Ludwig Maximilian Universität, München.

Christoffersen, Peter, Vihang Errunza, Kris Jacobs, and Hugues Langlois (2012). "Is the Potential for International Diversification Disappearing? A Dynamic Copula Approach." *Review of Financial Studies*, 25(12), 3711–3751.

Galati, Gabriele, Steven Poelhekke, and Chen Zhou (2011). "Did the Crisis Affect Inflation Expectations?" *International Journal of Central Banking*, 7(1), 167–207.

Gürkaynak, Refet, Andrew Levin, and Eric Swanson (2010). "Does inflation targetting anchor long-run inflation expectations? Evidence from the US, UK, and Sweden." *Journal of the European Economic Association*, 8(6), 1208–1242.

Nautz, Dieter and Till Strohsal (2015). "Are US inflation expectations re-anchored?" *Economics Letters*, 127, 6–9.

Nelsen, R. (2006). *An Introduction to Copulas*. Second ed., Springer, New York.

Oh, Dong Hwan and Andrew J. Patton (2013). "Time-Varying Systemic Risk: Evidence from a Dynamic Copula Model of CDS Spreads." Mimeo, Duke University.

Patton, Andrew J. (2004). "On the Out-of-Sample Importance of Skewness and Asymmetric Dependence for Asset Allocation." *Journal of Financial Econometrics*, 2(1), 130–168.

Patton, Andrew J. (2006a). "Estimation of Multivariate Models for Time Series of Possibly Different Lengths." *Journal of Applied Econometrics*, 21, 147–173.

Patton, Andrew J. (2006b). "Modelling asymmetric exchange rate dependence." *International Economic Review*, 47(2), 527–556.

Patton, Andrew J. (2012). "A review of copula models for economic time series." *Journal of Multivariate Analysis*, 110, 4–18.

Resnick, Sidney I. (2007). *Heavy-tail phenomena: probabilistic and statistical modeling*. Springer.

Sklar, A. (1959). "Fonctions de répartition à N dimensions et leurs marges." *Publ. Inst. Statist. Univ. Paris*, 8, 229–231.

Sklar, A. (1973). "Random variables, joint distributions, and copulas." *Kybernetica*, 9, 449–460.

Trivedi, P. K. and D. M. Zimmer (2005). "Copula Modeling: An Introduction for Practitioners." *Foundations and Trends in Econometrics*, 1(1), 1–111.

Appendix A: Tail dependence

In this article, attention is restricted to the bivariate case; in most instances the theoretical generalization to the m-dimensional case is straightforward. It is useful to provide some theoretical background. Given two random variables X and Y, define the joint cumulative distribution function F as:

$$F(x,y) = \Pr\{X \leq x \text{ and } Y \leq y\}\,. \tag{A.1}$$

In order for F to qualify as a cumulative distribution function, it has to fulfill certain requirements. Intuitively, it is clear that F has to be 0 if any of its arguments is below the lowest value that the respective random variable can attain; it has to be 1 if all its arguments are higher than the highest value that each random variable can attain; and it must assign a non-negative value for the probability of any rectangle in its domain. Formally, these ideas would be expressed as $\lim_{x\to-\infty} F(x,y) = 0$ (and similarly for y), $\lim_{x,y\to+\infty} F(x,y) = 1$, and $F(x_1,y_1) + F(x_2,y_2) - F(x_1,y_2) - F(x_2,y_1) \geq 0$ for any (x_1,y_1) and (x_2,y_2).

The one-dimensional margins are defined as $F_X(x) = \lim_{y\to+\infty} F(x,y)$ and $F_Y(y) = \lim_{x\to+\infty} F(x,y)$. Let x_k denote quantile k of variable X, that is, the value of x that solves equation $F_X(x) = k$, and similarly for y.[7] The *conditional upper tail dependence* is defined as

$$\lambda_U = \lim_{k\to1} \Pr\{y > y_k | x > x_k\}\,. \tag{A.2}$$

Similarly, it is possible to define the lower tail dependence λ_L taking the limit as k goes to zero and reversing the inequalities.

Appendix B: More about copulas

The first important characteristic of a copula is that its underlying random variables are defined in the $[0,1]$ interval. The second important characteristic is that the copula's marginal distributions are uniform. Copulas are relevant because they connect multivariate distributions to their one-dimensional margins. Under pretty standard regularity conditions, a theorem due to Sklar (1959, 1973) states that there exists a copula C satisfying $F(x,y) = C(F_X(x), F_Y(y))$. In other words, any bi-dimensional cumulative distribution function can be decomposed into its marginal distributions and a copula. Moreover, the latter completely characterizes the dependence between the two variables. If the marginal cumulative distribution functions are continuous, this copula is unique.

7. The conditional cumulative distribution functions are, in case F is differentiable, $F_{X|Y}(x,y) = \frac{\partial F}{\partial y}(x,y)$ and $F_{Y|X}(x,y) = \frac{\partial F}{\partial x}(x,y)$.

One important consequence of this is that using the inverse of the marginal cumulative distribution function of X, F_X^{-1}, to transform a uniformly distributed variable in $[0,1]$, U, yields a variable that is distributed according to F_X. The same happens for Y and a uniformly distributed variable V in $[0,1]$. Therefore, $(F_X(x), F_Y(y))$ has copula C as its cumulative distribution function, and $(F_X^{-1}(u), F_Y^{-1}(v))$ has F as its cumulative distribution function. Because order relations in equation (A.2) are maintained between (x, y) and the corresponding uniformly distributed values (u, v), conditional tail dependence occurring for F will also occur for C.

Copulas turn out to be especially useful because tail dependence can be easily computed from their functional forms. Moreover, their domain fits nicely to the language of quantiles and percentiles necessary to study co-movement. There are a few notable copulas, some of which will be used in the body of this article. See Trivedi and Zimmer (2005) and Nelsen (2006) for a thorough exploration of different copulas and their properties. It is enough here to give just one example, which will be the Gumbel copula. Its expression is

$$C(u,v) = \exp\left(- \left((-\log(u))^{\theta} + (-\log(v))^{\theta} \right)^{\frac{1}{\theta}} \right), \tag{B.1}$$

where $\theta \in [1, +\infty]$. If θ is 1, the copula collapses to $C(u,v) = uv$, which is the case where variables are independent. If θ goes to $+\infty$, then $C(u,v) = \min\{u, v\}$, which corresponds to maximum dependence; this would imply correlation 1 between the two variables. This copula does not exhibit lower tail dependence, which may or may not be an obstacle to its utilization, but in turn can display arbitrarily large upper tail dependence. If one is interested in focusing on the co-movement between large upward revisions of short-term inflation expectations and long-term inflation expectations, then a Gumbel copula would be appropriate.[8] The formula above also allows for the computation of the upper tail dependence as expressed by equation (A.2); the result is $\lambda_U = 2 - 2^{\frac{1}{\theta}}$. As θ approaches 1 upper tail dependence approaches 0, which means no dependence; as θ approaches $+\infty$ upper tail dependence approaches 1, which means full correlation between the upper tails of the two variables. Figure B.1 provides a visual representation of the Gumbel copula for several levels of tail dependence: θ equal to 1, 1.3, 2.5 and $+\infty$, which entail upper tail dependence of 0, 0.3, 0.68 and 1, respectively. Several of the typical characteristics of copulas are evident. First, the marginal distributions are uniform, as can be seen from the straight line segments connecting (1,0,0) to (1,0,1), and (0,1,0) to (0,1,1). Second, as tail dependence increases from the independence case ($\theta = 1$) to the full correlation case ($\theta \rightarrow \infty$) the iso-probability curves (the "level curves" in the copula graph) go from hyperbolas

8. In fact, it is also possible to study lower tail dependence using the so-called Rotated Gumbel copula, whose expression is that in (B.1) with the arguments replaced by $1 - u$ and $1 - v$.

(with equation $k = uv$) to two segments connected at right angles at points such that $u = v$.

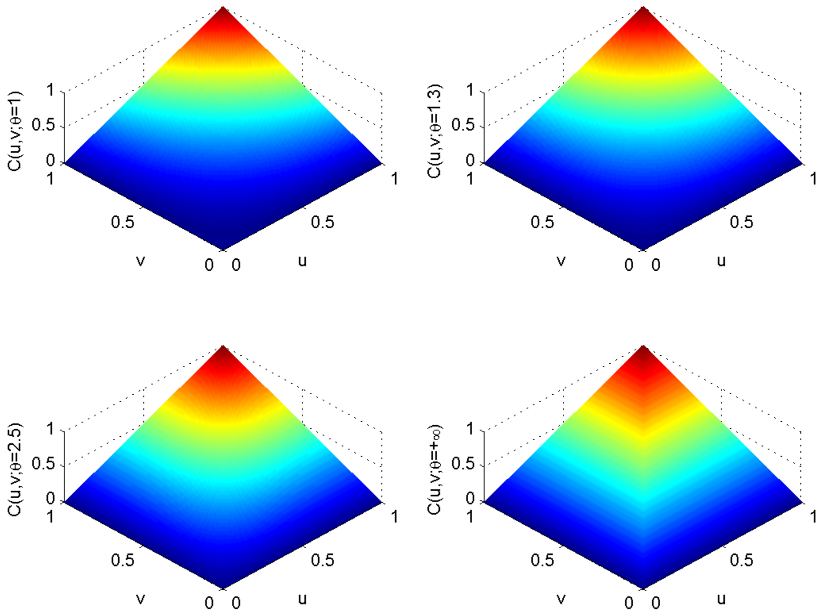

FIGURE B.1: Gumbel copula for several values of θ.

Source: author's calculations.

A last topic in terms of copulas concerns dynamic copulas. Dynamic copulas were first introduced by Patton (2006b) and are essentially the same as static copulas except that a subset of, or all, the parameters governing dependence is allowed to change over time. Patton (2006a), Braun and Grziska (2011) and Oh and Patton (2013) provide examples of dynamic copulas. The way in which parameters evolve over time is somehow arbitrary. Several dynamic copulas were also estimated for the data used here. The results do not differ significantly from those reported in this article and are available from the author upon request.