Visualizing the Choices: Marginals, Risk Measures & Event Frequency

Uncertainty and Risk Visualisation

KTN-IM, KTN-FS, TSB, NERC Workshop

Lloyds, May 2010

William Shaw
Professor of Financial Mathematics
King's College London
http://www.mth.kcl.ac.uk/staff/w_shaw.html
william.shaw@kcl.ac.uk

Goal: clarify importance of maths ingredients

Within the realm of risk computations addressable by mathematical and statistical analysis, we have some choices to make in quoting a result

Choice of consequence to report (temp increase, ... area flooded,... population homeless....)

Choice of measure to quote (VaR/quantile, CVaR/ETL,... other)

Choice of marginal distributions

Choice of dependency structure

Choice of event frequency

I will not pretend this is the whole story. I do not worry about probability of a fire before insuring my house. I do not ensure computers because premium is large fraction of replacement. But within the statistically treatable domain these choices are critical.

Do organizations create a thin-tailed bias?

Allow me a little mischief:

Survival within organizations may require (a) budget control (do not waste money on the improbable) and (b) ready scapegoat for problems.

Hilary Benn, the Environment Secretary, on the Nov 2009 floods: whilst the area's flood defences had been built to withstand a "one in 100 year" flood, "*what we dealt with last night was probably more like one in a thousand, so even the very best defences, if you have such quantities of rain in such a short space of time, can be over-topped*".

David Viniar, CFO, Goldman Sachs:
We were seeing things that were 25-standard deviation moves, several days in a row.

Interesting figure in relation to Gaussian modelling, in which a single 25σ event would not be expected in one lifetime of the universe. You need about 10^{15} *Googol* universe lifetimes!

Who or what do we blame?

During the recent crisis there were various whinges and excuses and blame games. The Mathematics community had a bit of a job on its hands to avoid being assigned the blame when the root cause of the problem was a simple failure in lending mortgage money to people who could not pay it back.

WIRED Mag story: "The formula that destroyed Wall Street" saying that it was all the fault of the Gaussian copula - which is hopelessly wrong but a long story, maybe at the end....

I mentioned Viniar and Benn. I tend to lump all these together as rewordings of CRO B Simpson

It wasn't my fault.
I didn't do it.
Nobody saw me do it,
you can't prove anything.

The general idea for "Risk Manager Simpson" is to claim a running assumption that the distribution of events is narrow, so that extreme events are very unlikely.

Some numbers

This is a live *Mathematica* V7 notebook that I am happy to distribute

In[1]:=
```
F[x_] := CDF[NormalDistribution[0, 1], x]

Plot[F[x], {x, -3, 3}]
```

Out[2]=

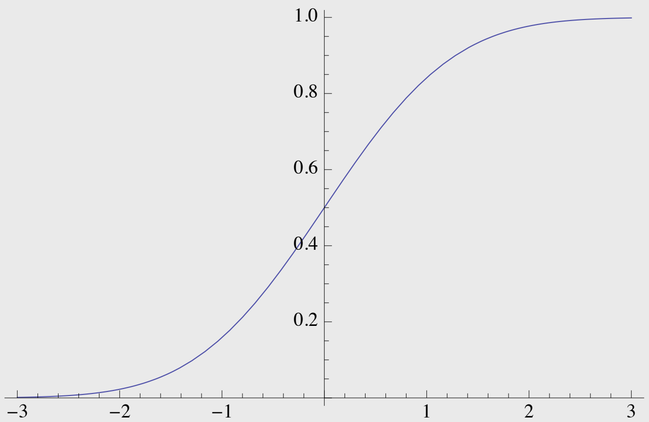

In[3]:=
```
F[-25] // N
```

Out[3]= 3.0567×10^{-138}

Reality, at least from finance

The assumption of normality that has pervaded risk computations has been known to be incorrect for decades, and it has been quantified in more detail well before the current crisis.

1960s - Mandelbrot and Fama describe non-Gaussian nature of asset (log-)returns

1970s - Black and Scholes Gaussian(!) model

Various studies....

2003 - MIT-BU group claim "universal" inverse *cubic* behaviour in CDF tail (*Nature*)

2006 - Platen team report MLE estimate of T_4 in major world indices (*Applied Math Finance*, MLE amongst generalized hyperbolic family), i.e. inverse *quartic* in CDF tail.

Likewise, Platen EPJ-B, Taylor S. Africa data (T, VG). These differ in the detail but agree data is massively non-normal - indeed super-fat with *possibly infinite model kurtosis*. (Try working out fourth moment for a pdf with density decaying like x^{-4}, x^{-5}).

Trying to do better

We ran a little exercise with PhD students from KCL, Imperial, Birkbeck and LSE over Xmas, where we looked at the ingredients for reporting risk numbers. We looked at 3 of the ingredients

Choice of risk measure (VaR=quantile, CVaR etc)
Choice of distribution
Choice of frequency level to report

The critical thing is the quantile or other risk measure, expressed here as how many std devs for a given frequency. So here is a bit of computational maths to set it up. We will in fact be ultra conservative and assume that there are (rather unrealistically)

NO shifts in the mean
NO explosion in the variance

so it is purely about the choice of distribution with zero mean and a fixed variance. Make good use of advanced math computation environment *Mathematica* V7 - advanced functions and easy app development.

Model Choice

I will use T-statistics (Gosset aka Student 1908). Not just suggested by some significant data statistics. There is some mathematical underpinning as well.

In finance if $X_t = \log(S_t / S_0)$, the Black-Scholes (1973) world is described by the SDE

$$\mathrm{dX}_t = \mu\, dt + \sigma\, dW_t$$

and log-returns are forever Gaussian.

If we introduce simple linear price feedback, possibly motivated by technical trading:

$$\mathrm{dX}_t = (\mu - \lambda\, X_t)\, dt + \sigma_1\, dW_{1\,t} + \sigma_2\, X_t\, dW_{2\,t}$$

then if the parameters are right (strong enough mean reversion) then equilibria of Student and Pearson IV type emerge, and are one of several dynamic outcomes. Others include a momentum-dominated panic. (Shaw 2009, builds on Nagahara 1996, Wong 1953!)

Even if you do not believe this model, the following hopefully illuminates what matters in the choices.The T_4 is multiply-special: closed-form quantile, found as MLE in index data, boundary of finite kurtosis.

Some Maths: Quantiles, VaR, CVaR formulae

Normal Quantile

In[4]:=
```
QuantileN[u_] := Sqrt[2] InverseErf[2 u - 1]
```

T_4 case: Shaw (JCF 2006) closed form

In[5]:=
```
InverseCDF4[y_] := Module[{ra = Sqrt[1 - 4 (y - 1 / 2) ^2]},
  2 Sign[y - 1 / 2] * Sqrt[Cos[1 / 3 ArcCos[ra]] / ra - 1]]
```

With all these power law VaR, we need to normalize to unit variance, assuming $n > 2$.

In[6]:=
```
Stunorm4[u_] := 1 / Sqrt[2] InverseCDF4[u]
```

T_n and T_3

In[7]:=
```
FMinusOne[y_, n_] := Module[{arg = If[y < 1 / 2, 2 y, 2 (1 - y)]},
  Sign[y - 1 / 2] Sqrt[n * (1 / InverseBetaRegularized[arg, n / 2, 1 / 2] - 1)]]
```

In[8]:=
```
Stunorm[u_, n_] := Sqrt[(n - 2) / n] FMinusOne[u, n]
```

In[9]:=
```
Stunorm3[u_] := 1 / Sqrt[3] FMinusOne[u, 3]
```

CVaR variations

To treat the expected tail loss, or conditional "Value at Risk", we compute the expected value GIVEN that one is in the tail. *In these cases the integrals can be done in closed form in terms of "density-like" functions evaluated at the quantile.*

Normal CVaR/ETL

In[10]:=
```
f[x_] := Exp[-x^2/2]/Sqrt[2 Pi]; NCVaR[u_] := -f[QuantileN[u]]/u
```

T meta-density

In[11]:=
```
metaf[t_, n_] := -(n^(n/2) * (n + t^2)^(1/2 - n/2) * Gamma[(-1 + n)/2]) / (2 * Sqrt[Pi] * Gamma[n/2])
```

T_4 CVaR/ETL

In[12]:=
```
StuCVaR4[u_] := 1/Sqrt[2] metaf[InverseCDF4[u], 4]/u
```

T_n CVaR/ETL

In[13]:=
```
StuCVaR[u_, n_] := Sqrt[(n - 2)/n] metaf[FMinusOne[u, n], n]/u
```

VaR values

One problem is banks reporting risk based only on high frequency events. Look at the 2.5% level (i.e. things will be this bad or worse 2.5% of the time, or once every 40 days). For contrast we will have some lower frequency events as well:

In[14]:=
```
uvals = {2.5 / 100, 1 / 100, 1 / 10000, 1 / 10^6}
```

Out[14]= $\left\{0.025, \frac{1}{100}, \frac{1}{10\,000}, \frac{1}{1\,000\,000}\right\}$

In the normal case the first number is the famous (-)1.96 figure:

In[15]:=
```
Map[QuantileN, uvals] // N
```

Out[15]= {−1.95996, −2.32635, −3.71902, −4.75342}

Note that making a switch to CVaR does NOT change matters much!

In[16]:=
```
Map[NCVaR, uvals] // N
```

Out[16]= {−2.3378, −2.66521, −3.95848, −4.94833}

Redo the normal case: the first number is the famous (-)1.96 figure:

In[17]:=
```
Map[QuantileN, uvals] // N
```
Out[17]= {−1.95996, −2.32635, −3.71902, −4.75342}

But switching the distribution is more interesting. ...

In[18]:=
```
Map[Stunorm4, uvals] // N
```
Out[18]= {−1.96324, −2.64949, −9.2162, −29.4}

In[19]:=
```
Map[Stunorm3, uvals] // N
```
Out[19]= {−1.83739, −2.62158, −12.8193, −59.64}

The once in 40 days number is about -2 no matter what model you use. This is what I call the maths behind the observation that "*Gaussian VaR is an airbag that always works until you need it*." But look at the other numbers, which go through the roof - the choice of tail model seriously matters.

So we need to get a better grip on all this, so here is the prototype FatRisk Tool....

Fat Tail Risk Visualization Tool (V1, VaR/Quantile)

In[20]:=
```
bennone = -Log[10, 1 / (260)] // N;
bennten = -Log[10, 1 / (10 * 260)] // N;
bennhundred = -Log[10, 1 / (100 * 260)] // N;
bennthousand = -Log[10, 1 / (1000 * 260)] // N;
Manipulate[onelevel = -Stunorm[1 / 260, n]; tenlevel = -Stunorm[1 / 260 / 10, n];
 hundredlevel = -Stunorm[1 / 260 / 100, n]; thousandlevel = -Stunorm[1 / 260 / 1000, n];
 Plot[{-QuantileN[10^(-x)], -Stunorm4[10^(-x)], -Stunorm[10^(-x), n]} , {x, 1, 6}, PlotRange → {0, 30},
  PlotLabel → Style["Daily VaR: Gauss, T_4, T_n (μ=0,σ=1)", "Large"], PlotStyle → Thickness[0.005],
  FrameLabel → {Style["-Log[10, freq]", Large], Style["Std. Devs.", Large]}, Frame → True, ImageSize → 600,
  Epilog → {{Line[{{1, 25}, {6, 25}}], Line[{{bennone, 0}, {bennone, 30}}], Line[{{bennten, 0}, {bennten, 30}}],
     Line[{{bennhundred, 0}, {bennhundred, 30}}], Line[{{bennthousand, 0}, {bennthousand, 30}}],
     Text["1Y", {2.42, 28}, Left], Text["10Y", {3.42, 28}, Left], Text["100Y", {4.42, 28}, Left],
     Text["1000Y", {5.42, 28}, Left], Text["10%", {1, 0.7}], Text["1%", {2, 0.7}], Text["0.1%", {3, 0.7}],
     Text["0.01%", {4, 0.7}], Text["0.001%", {5, 0.7}], Text["0.0001%", {5.95, 0.7}]},
    PointSize[0.02], Point[{6 - Log[10, 1.91], -Stunorm4[1.91 × 10^(-6)]}],
    Text[Style["OK airbag", Larger], {1.6, 5}], Text["Viniar level", {1.5, 25.5}],
    Line[{{0.5, onelevel}, {bennone, onelevel}}],
    Line[{{0.5, tenlevel}, {bennten, tenlevel}}],
    Line[{{0.5, hundredlevel}, {bennhundred, hundredlevel}}],
    Line[{{0.5, thousandlevel}, {bennthousand, thousandlevel}}],
    {PointSize[0.01], Point[{bennone, onelevel}], Point[{bennten, tenlevel}],
     Point[{bennhundred, hundredlevel}], Point[{bennthousand, thousandlevel}]}}], Delimiter,
 Dynamic[Panel[TableForm[{{"Power", n}, {"1Y VaR", N[onelevel]}, {"10Y VaR", N[tenlevel]},
     {"100Y VaR", N[hundredlevel]}, {"1000Y VaR", N[thousandlevel]}}]]],
 {{n, 4, "CDF Power Law"}, 10, 2.1, 0.001}, {{n, 4, "CDF Power Law"}, {2.25, 2.5, 3, 4, 5, 6, 7, 8}},
 ContinuousAction → True, ContentSize → 650, ControlPlacement → Bottom]
```

Out[24]=

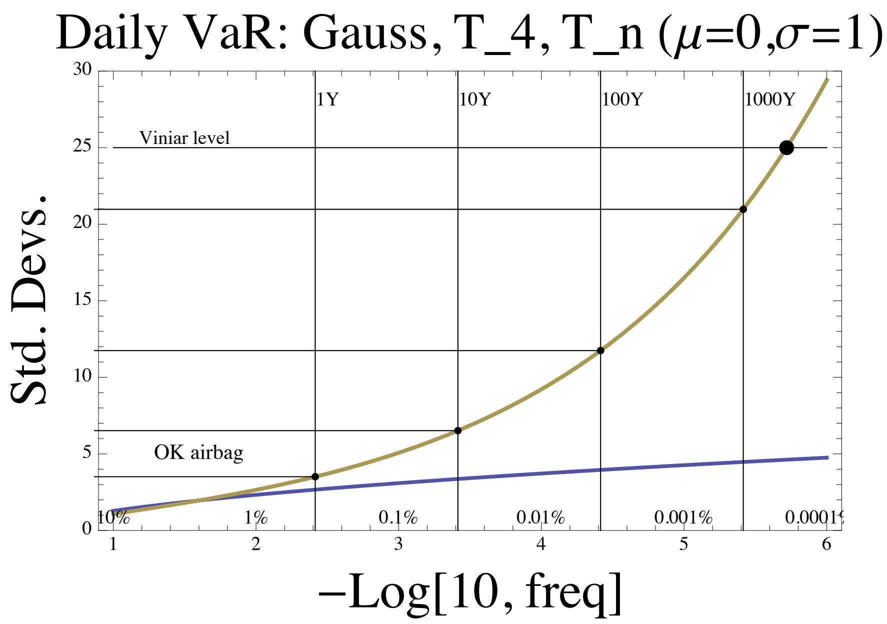

Power	4
1Y VaR	3.50903
10Y VaR	6.51897
100Y VaR	11.7464
1000Y VaR	20.9743

CDF Power Law

CDF Power Law 4

To Consider

Proximity of curves in low frequency range. The Gaussian airbag is OK - you would barely notice you were using the wrong distribution based on fortnight to annual comparisons

Let's look at T_4, suggested by Fergusson-Platen (2006) study. 10Y VaR is 6.5sd, still not achieved in Gaussian for 1 in a million day event.

If we believe the MIT-BU analysis, n=3, then we are more extreme territory still and an 8 sigma event happens once a decade on average.

power = 2 singular limit as std dev heads off to infinity and we cannot normalize in this way. Can explore extreme and particular values with ease.

You need to be clear on the distribution and frequency choices to have any clue as to the possibilities.

Related Issues

There are all kinds of other problems. The choice of risk number to report, and the choice of dependency model are critical too. The latter needs to captures the domino effect better.

The use of VaR has been widely criticised in the finance community. Given a number, say 1%, this is the quantile at 0.01.

There are presentational issues linked to this. Which sounds better?

"The loss is better than this 99% of the time"
"Things will be this bad or worse 1% of the time"

Interestingly, the climate change work is sometimes using this same measure, just as finance has got their head around dumping it.

A mathematically better idea is to use the expected loss for the worst 1% of occasions - aggregates better and captures tail details. But how much difference does it REALLY make?

How does use of CVaR/expected tail loss modulate?

```
Manipulate[Plot[{-NCVaR[10^(-x)], -StuCVaR4[10^(-x)], -QuantileN[10^(-x)],
   -Stunorm4[10^(-x)], -Stunorm[10^(-x), n], -StuCVaR[10^(-x), n]} , {x, 1, 6},
  PlotRange → {0, 30}, PlotLabel → Style["Daily CVaR vs VaR: \n Gaussian[B] vs T_4[G], T_3[B]", "Large"],
  PlotStyle → {{Thickness[0.0025], Blue}, {Thickness[0.0025], Green}, {Thickness[0.00125], Blue},
    {Thickness[0.00125], Green}, {Thickness[0.00125], Red}, {Thickness[0.0025], Red}},
  FrameLabel → {Style["-Log[10, freq]", Large], Style["Std. Devs.", Large]}, Frame → True,
  ImageSize → 700, Epilog → {{Text["10%", {1, 0.7}], Text["1%", {2, 0.7}], Text["0.1%", {3, 0.7}],
     Text["0.01%", {4, 0.7}], Text["0.001%", {5, 0.7}], Text["0.0001%", {5.95, 0.7}]},
    Line[{{bennone, 0}, {bennone, 30}}], Line[{{bennten, 0}, {bennten, 30}}],
    Line[{{bennhundred, 0}, {bennhundred, 30}}], Line[{{bennthousand, 0}, {bennthousand, 30}}],
    Text["1Y", {2.42, 28}, Left], Text["10Y", {3.42, 28}, Left], Text["100Y", {4.42, 28}, Left],
    Text["1000Y", {5.42, 28}, Left]}], {{n, 4, "CDF Power Law"}, 10, 2.1, 0.001},
 {{n, 4, "CDF Power Law"}, {2.25, 2.5, 3, 4, 5, 6, 7, 8}}, ContinuousAction → True,
 ContentSize → 1000, ControlPlacement → Bottom]
```

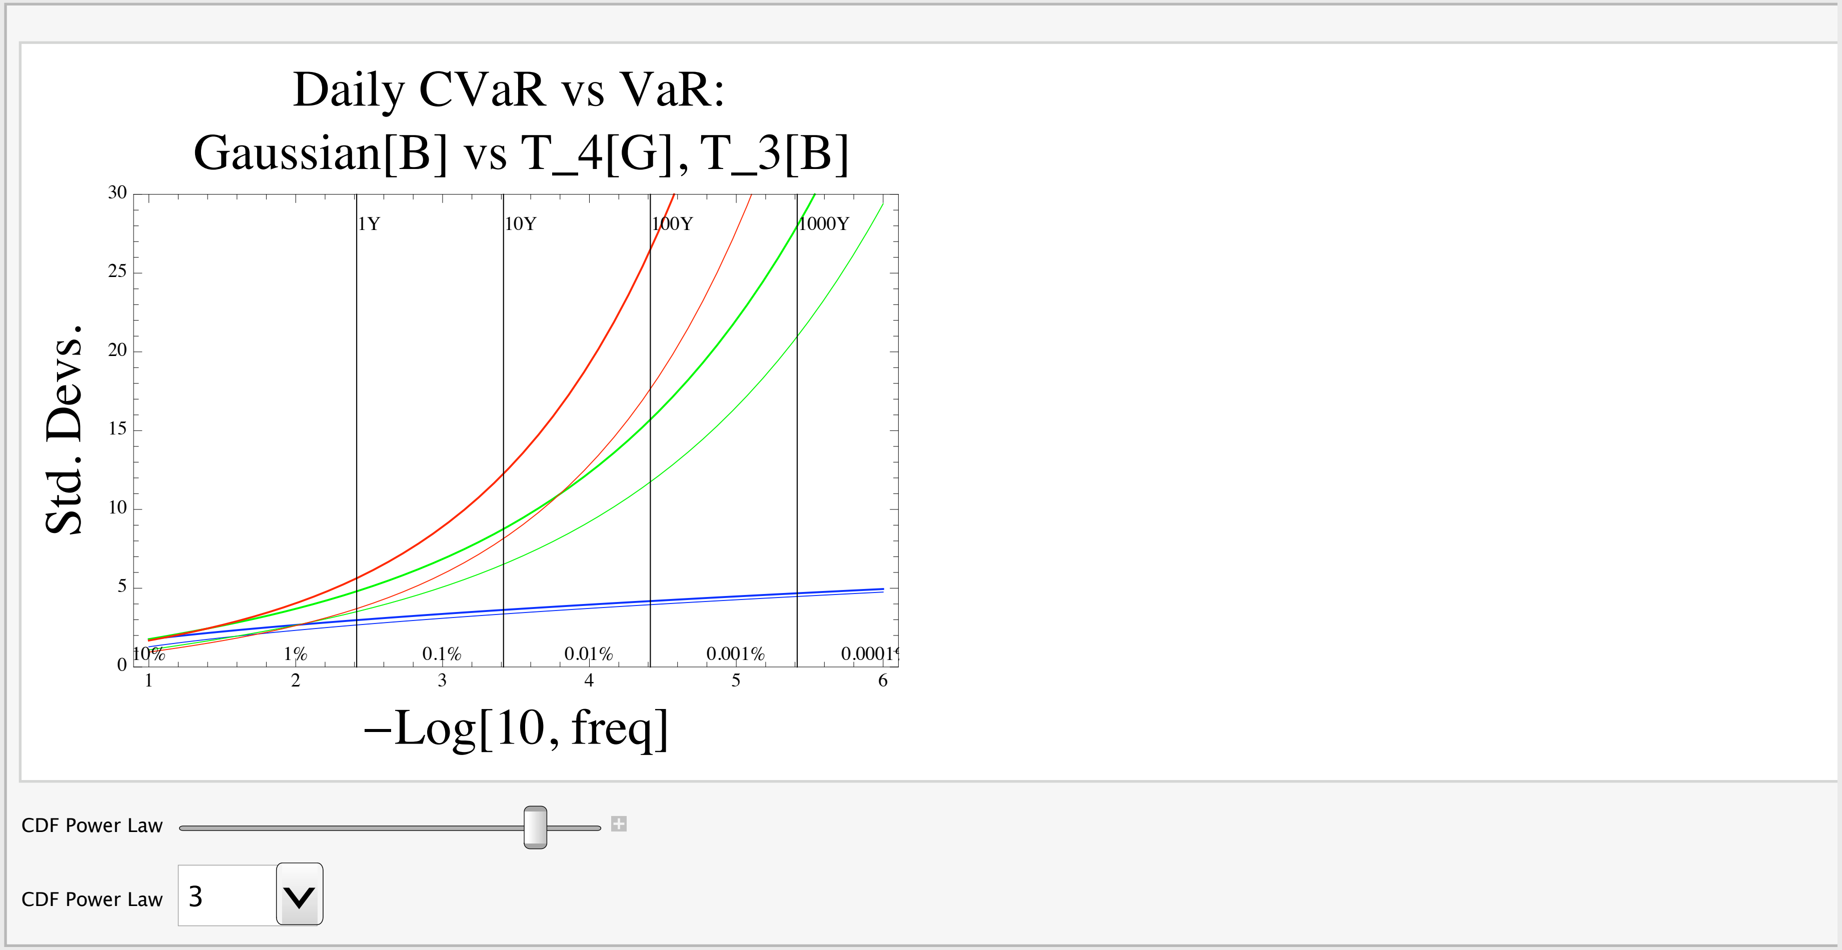

It is an almost irrelevant switch if you remain Gaussian. Choice of model is critical.

Dependency - quick look

This talk has mainly been about marginals. However, the dependency structure matters too. There are many types of dependency:

Company A has a real influence on company B
Both A, B influenced by common external factor
Spurious associations (same sector)

It is clear that a correlation number does not capture all the possibilities. This is not an excuse for throwing rocks at those who tried to come up with models to couple systems in a tractable manner.

A low point in the discussion of the crisis has to be Felix Salmon's hype in Wired Magazine. Despite getting quotes from Wilmott and Taleb that, when read carefully, made it pretty clear that the real problem was a naive (constant, historical, poor) choice of correlation, Salmon tried to pin a lot of blame on the Gaussian copula.

Exercise: Sample from Gauss, T (pick your own dof) Clayton copulas with Kendall's τ varying and measure your favourite risk with function the same marginals. Deduce that the Gaussian copula is not the problem!

Here is the initialization code (RUN IT)

In[25]:=
```
ClaytonParamFromTau[KendallTau_] := Module[{τ = KendallTau}, 2 τ / (1 - τ)]
```

In[26]:=
```
GaussCorrFromTau[KendallTau_] := Sin[Pi * KendallTau / 2]
```

In[27]:=

$$\texttt{Ncdf[x\_] :=} \frac{1}{2}\left(\texttt{Erf}\left[\frac{x}{\sqrt{2}}\right] + 1\right)$$

In[28]:=
```
uncorrsamp = RandomReal[NormalDistribution[0, 1], {1000, 2}];
```

In[29]:=
```
MakeGaussGraph[τ_] :=
 (ρ = GaussCorrFromTau[τ] ;
  L = {{1, 0}, {ρ, Sqrt[1 - ρ^2]}};
  corrsamp = (L.#1 &) /@ uncorrsamp;
  copsamp = (Ncdf[#1] &) /@ corrsamp;
  ListPlot[copsamp, AspectRatio → 1, PlotStyle → PointSize[0.005], PlotLabel → "Gaussian"])
```

In[30]:=
```
Clear[a, m, x, k]
```

In[31]:=
```
a[0, m_] := Gamma[(m + 1) / 2] / Gamma[m / 2] / Sqrt[m Pi];
a[k_, m_] := a[k, m] = (m - 2 k) / m / (2 k + 1) a[k - 1, m];
```

In[33]:=
```
TCDF[m_, x_] :=
 1 / 2 + x * Sum[a[p, m] x^(2 p), {p, 0, m / 2 - 1}] / (1 + x^2 / m)^(1 / 2 (m - 1))
```

In[34]:=
```
TCDF[2, x]
```

Out[34]=
$$\frac{x}{2\sqrt{2}\sqrt{\frac{x^2}{2}+1}}+\frac{1}{2}$$

In[35]:=
```
tdenoms = Sqrt[1 / 2 RandomReal[ChiSquareDistribution[2], {1000}]];
```

In[36]:=
```
tuncorrsamp = uncorrsamp / tdenoms;
```

In[37]:=

```
MakeTGraph[t_] := (ρt = GaussCorrFromTau[t] ;
L = {{1, 0}, {ρt, Sqrt[1 - ρt^2]}};
tcorrsamp = (L.#1 &) /@ tuncorrsamp;
tcopsamp = (TCDF[2, #1] &) /@ tcorrsamp;
ListPlot[tcopsamp, AspectRatio → 1, PlotStyle → PointSize[0.005], PlotLabel → "T_2 copula"])
```

In[38]:=

```
Claytoncorrpair[δ_] :=
  Module[{vone = RandomReal[], vtwo = RandomReal[], X = Random[GammaDistribution[1/δ, 1]]},
   {(1 - Log[vone]/X)^(-1/δ), (1 - Log[vtwo]/X)^(-1/δ)}]
```

In[39]:=

In[40]:=

```
MakeClayGraph[ct_] := (clay = Table[Claytoncorrpair[ClaytonParamFromTau[ct]], {1000}];
clayplotsmall = ListPlot[clay, AspectRatio → 1, PlotStyle → PointSize[0.005], PlotLabel → "Clayton"])
```

Correlation spiking against copula choice

The down tail starts to look the same once you spike the rank correlation measure - here Kendall's τ:

```
Manipulate[GraphicsArray[{MakeGaussGraph[τ], MakeTGraph[τ], MakeClayGraph[τ] }, ImageSize → 800],
 {{τ, 0.5, "Kendall τ"}, 0.001, 0.97}, {{τ, 0.5, "Kendall τ"}, N[Range[1, 9] / 10]}, ContinuousAction → True]
```

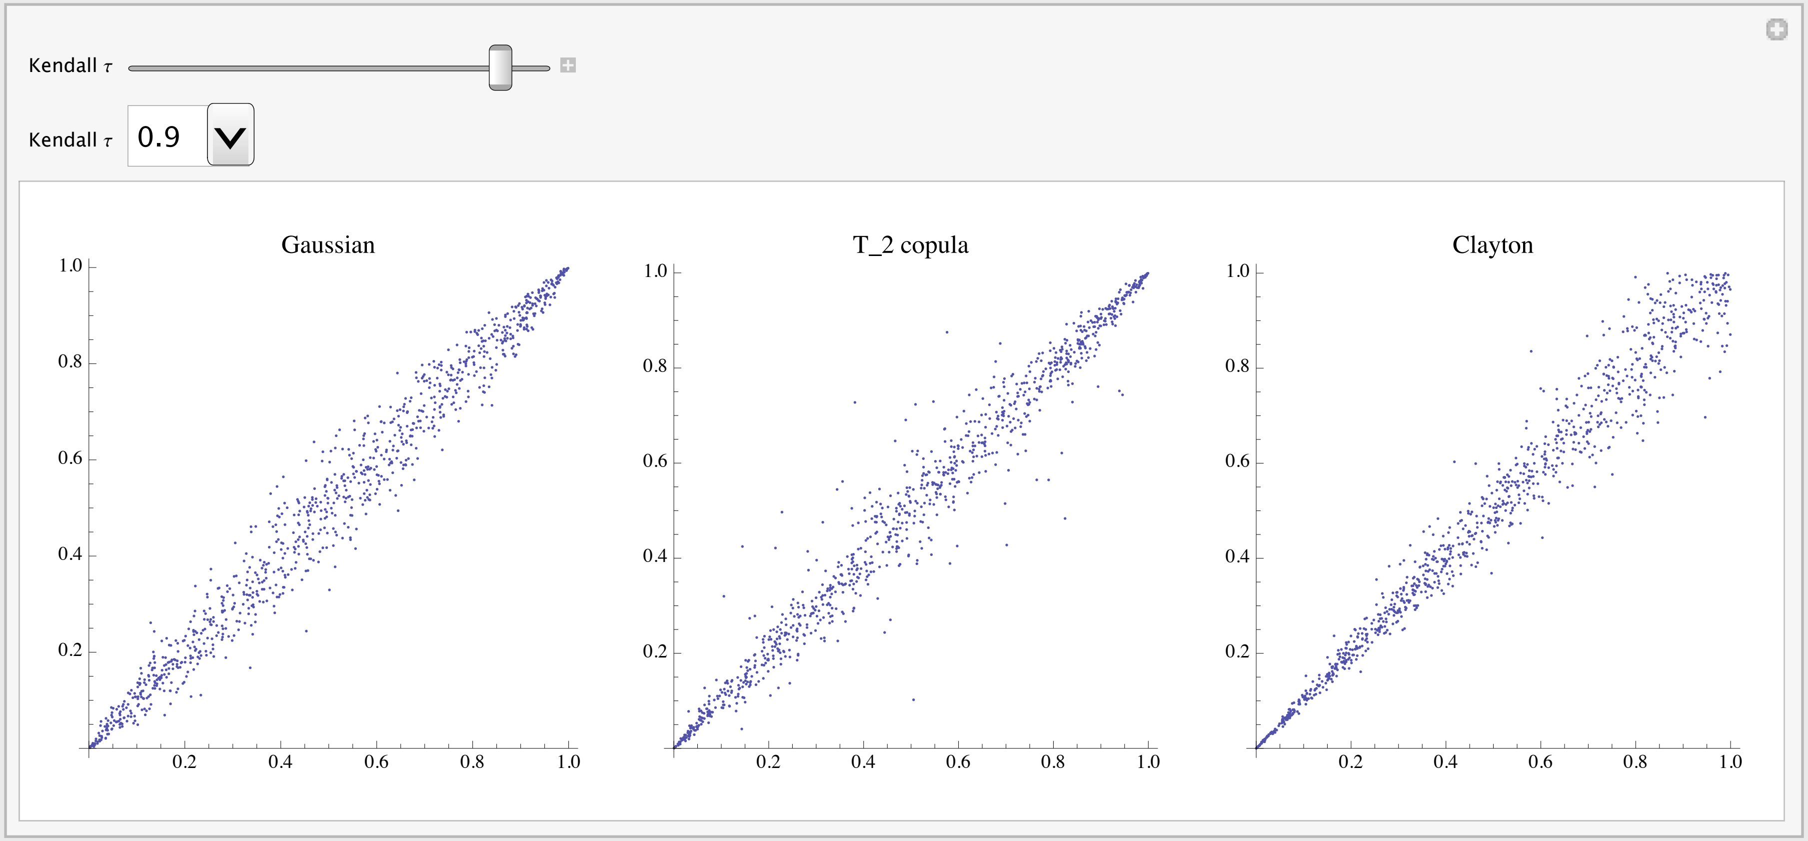

Kendall τ
Kendall τ 0.9
Gaussian
T_2 copula
Clayton
0.2 0.4 0.6 0.8 1.0

```
Manipulate[GraphicsArray[{MakeGaussGraph[τ], MakeTGraph[τ] }, ImageSize → 800],
 {{τ, 0.5, "Kendall τ"}, -0.97, 0.97}, {{τ, 0.5, "Kendall τ"}, N[Range[1, 9] / 10]}, ContinuousAction → True]
```

Kendall τ

Kendall τ

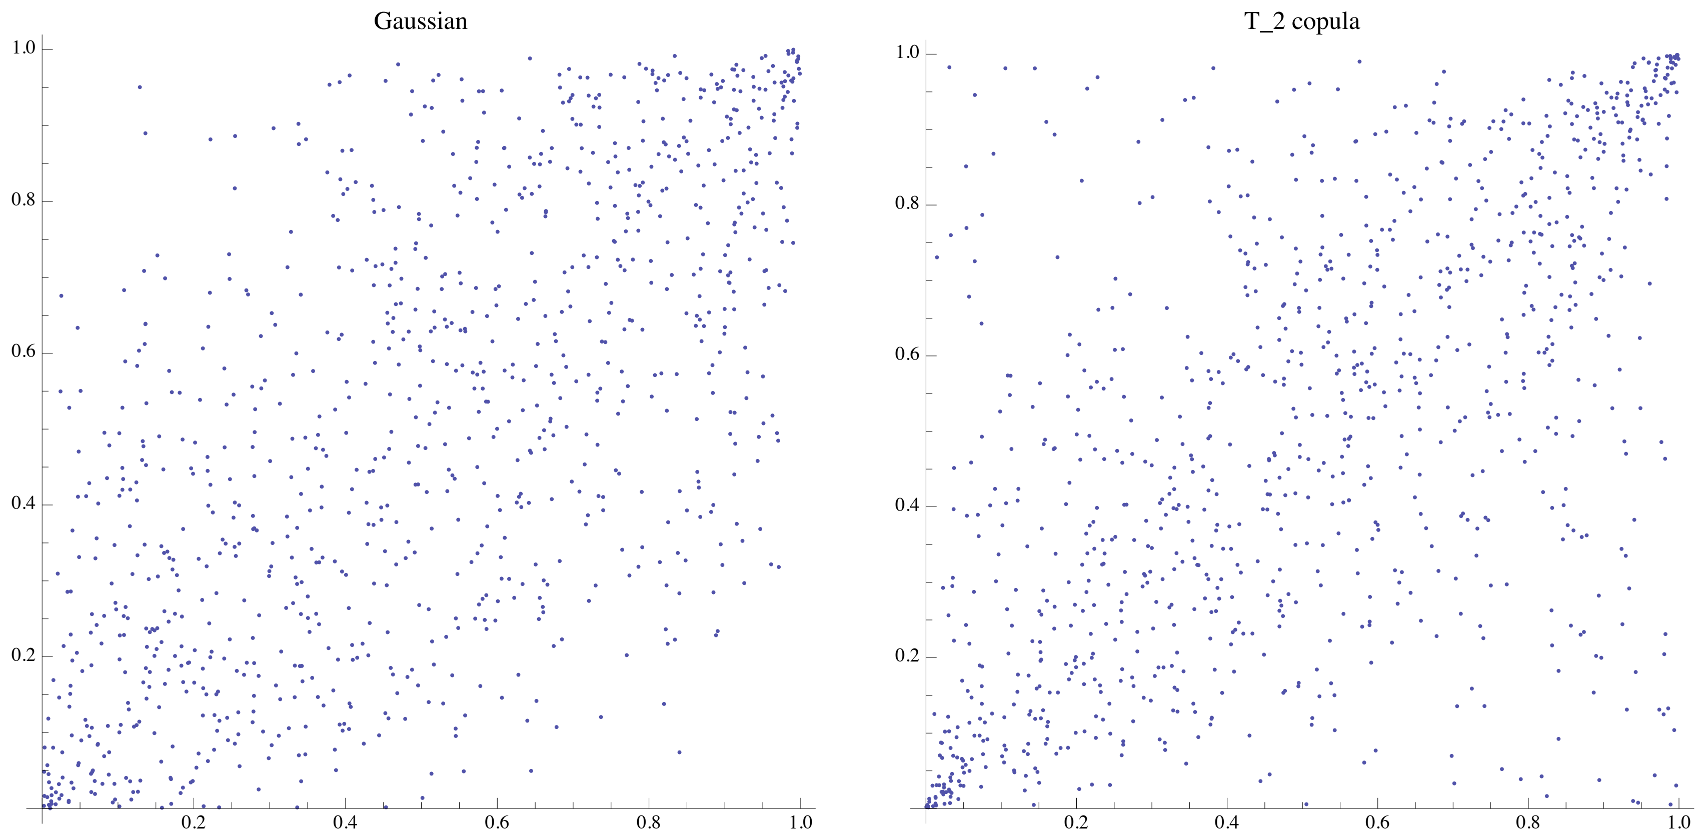

Summary

The choice of consequence, risk measure, event frequency and underlying marginal distribution are critical.

Switching VaR/quantile to CVaR/ETL almost irrelevant if you stay Gaussian

Switching Gaussian to pragmatic choice of T_n makes dramatic difference to lower frequency VaR values, but little difference around 2.5% level!

We can see what is going on easily with visualization tools like these, especially in looking at scale of 1,10, 100, 1000Y events in various T-models and comparing with Gaussian, in VaR or CVaR (other) risk measures.

Dependency remains complex, but even the simplest of tools can still give insight if properly used.

Interesting directions in risk optimization combining general marginals, dependencies, risk measures and robustness.