Why do you still use parametric distributions in empirical finance?

Enrique Sentana

CEMFI

XXXI Simposio de Análisis Económico

Oviedo, December 14^{th}, 2006

Squaring the circle

This lecture is supposed to be:

1. Relevant
2. Interesting to economists at large
3. Non-technical
4. Thought-provoking
5. Controversial
6. Funny
7. Memorable

Squaring the circle

This lecture is supposed to be:

1. Relevant
2. Interesting to economists at large
3. Non-technical
4. Thought-provoking
5. Controversial
6. Funny
7. Memorable

I'll do my best, but don't hold your breath

You should've known better!


The PRESIDENT

Motivation

- Nowadays, economic or finance theorists know more parametric distribution theory than either theoretical or empirical econometricians, unless they are Bayesian.
- A typical graduate student in a top PhD programme knows the univariate normal distribution, possibly the multivariate one, and the chi-square distribution, mostly because of their role in asymptotic theory.
- This is even worse than when I was a undergraduate student twenty years ago, because we also knew the t and F distributions.
- But these two distributions have long been forgotten because nobody believes any longer in the assumptions underlying the classical regression model.
- In fact, these days one doesn't use OLS in public any more. Instead, one applies GMM to the normal equations.

Motivation

- In contrast, economic theorists, especially those working in information economics or learning, know the gamma distribution, the beta distribution, and even the generalised inverse Gaussian distribution.
- Similarly, many applied statisticians, including actuarial scientists and those working on quality control, marketing, or credit scoring, regularly use parametric models based on those distributions.
- Moreover, they unashamedly use maximum likelihood methods to estimate their parameters.

Why is this?

Motivation

- Certainly not because of computational issues:
  1. From a numerical point of view, likelihoods are usually better behaved than generalised minimum distance objective functions.
  2. Besides, our laptops are far more powerful than the mainframes the Apollo programme used to put the first man on the Moon.
- The main reason is that we want our empirical analysis to be **'robust to misspecification'** at any cost, because economic theory is usually silent about distributional characteristics.
- As a result, we use either semiparametric methods such as GMM when we are interested in specific features of a distribution, or fully non-parametric ones when we are interested in a functional.

But is this really the right thing to do? I'm not sure

Motivation

- Robust to influential observations is another sense of the same word, but which is often incompatible with the previous one because sample means are particularly sentitive to outliers.
- The usual asymptotic theory of semiparametric and non-parametric estimators is often unreliable in finite samples.
- In fact, there are circumstances in which it simply doesn't apply because the required regularity conditions don't hold.
- In any case, given that our datasets are finite, I don't know that we should have lexicographic preferences over consistent estimators regardless of how inefficient they may be.

Motivation

Sampling distributions: Inconsistent vs. consistent estimator

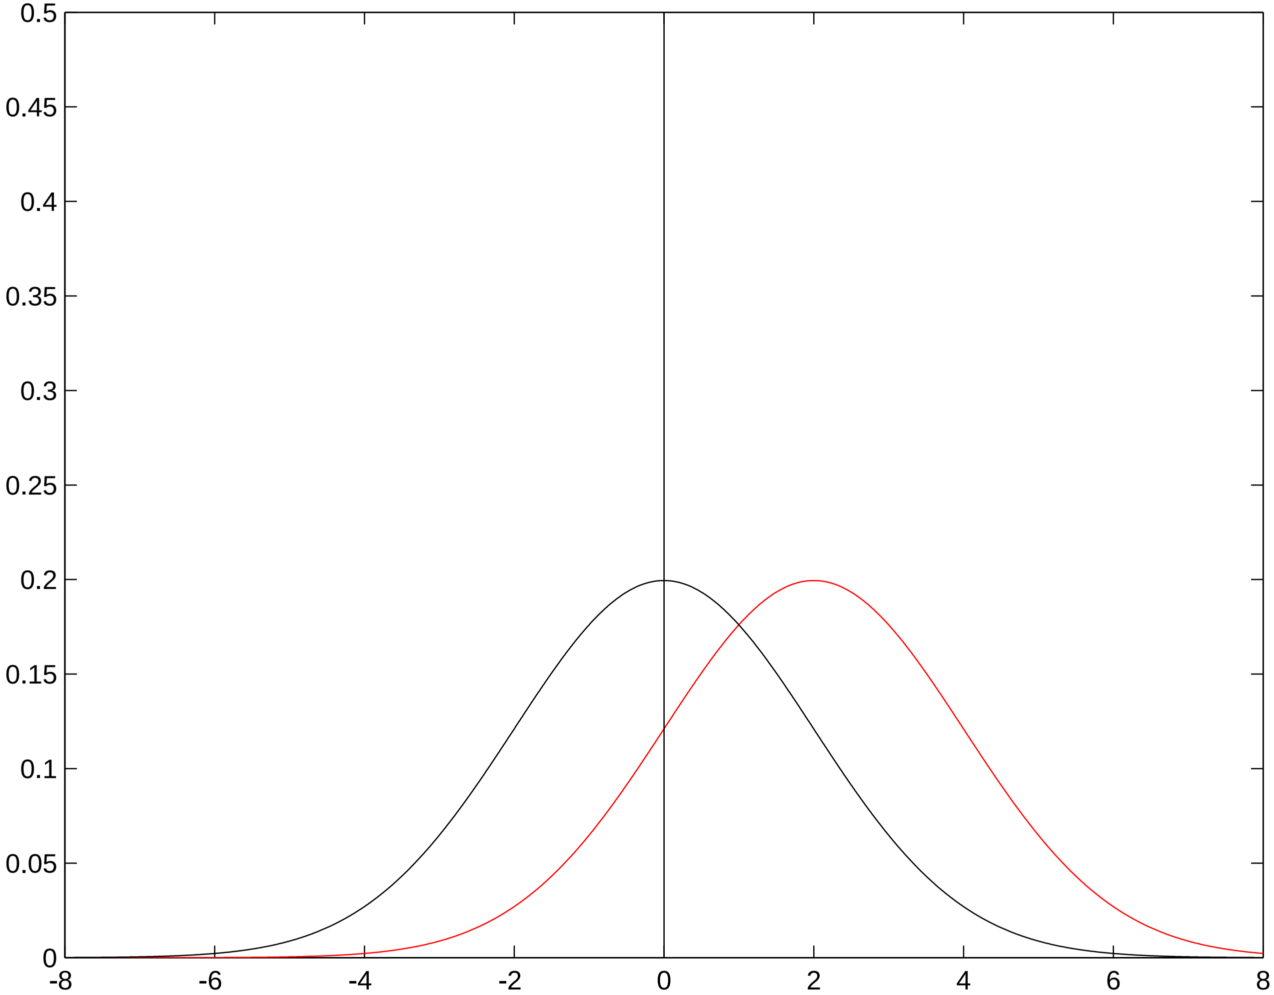

Motivation

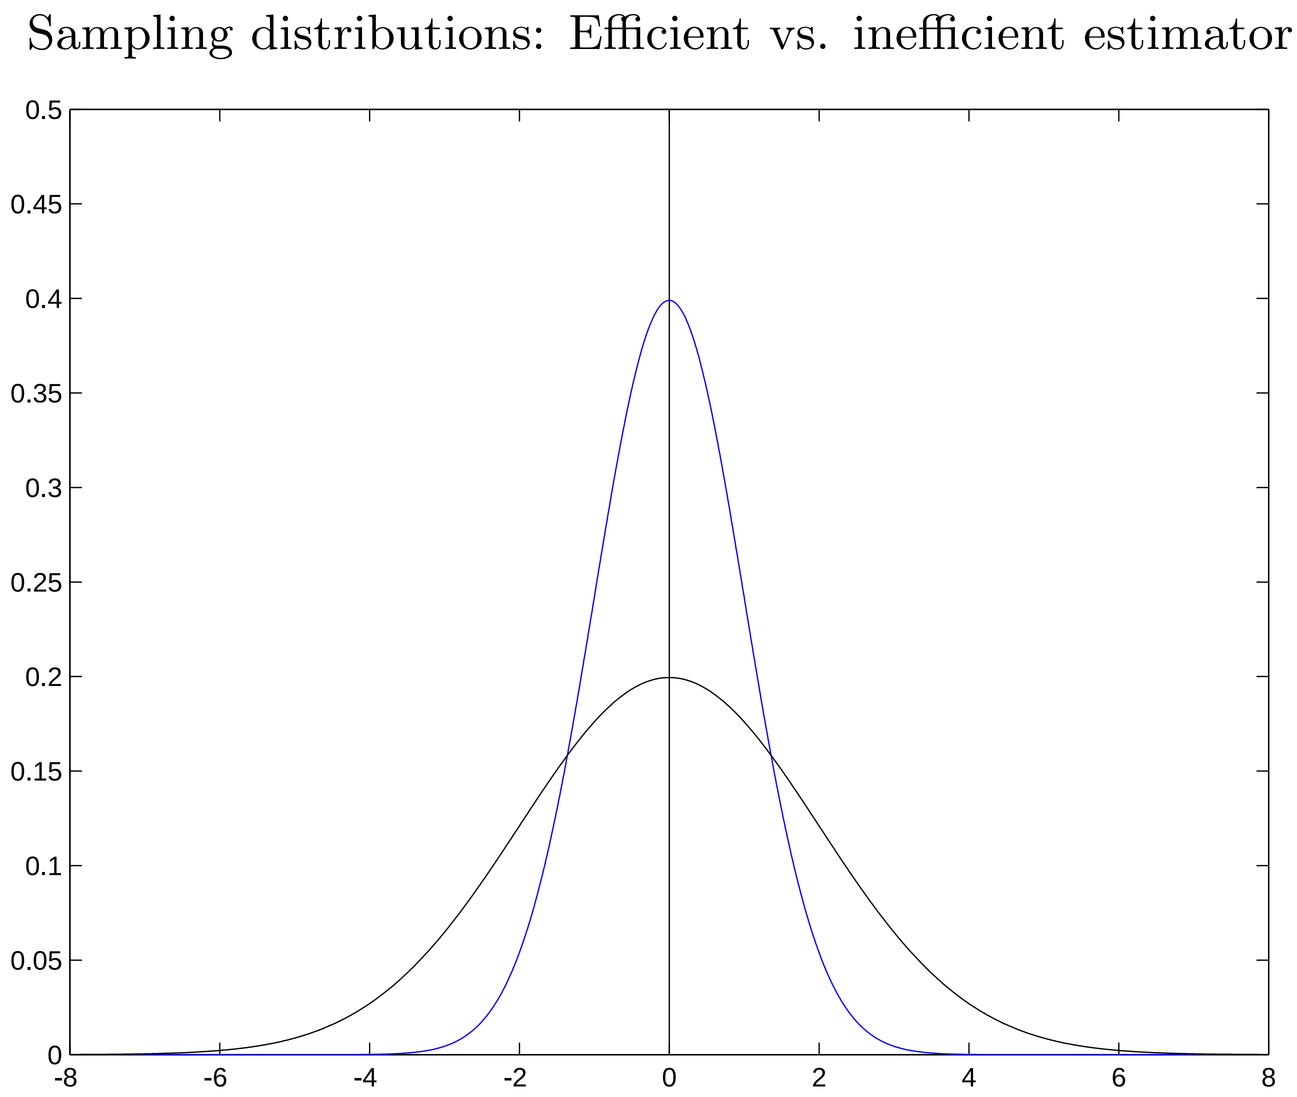

Sampling distributions: Efficient vs. inefficient estimator
0.5
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
-8
-6
-4
-2
0
2
4
6
8

Motivation

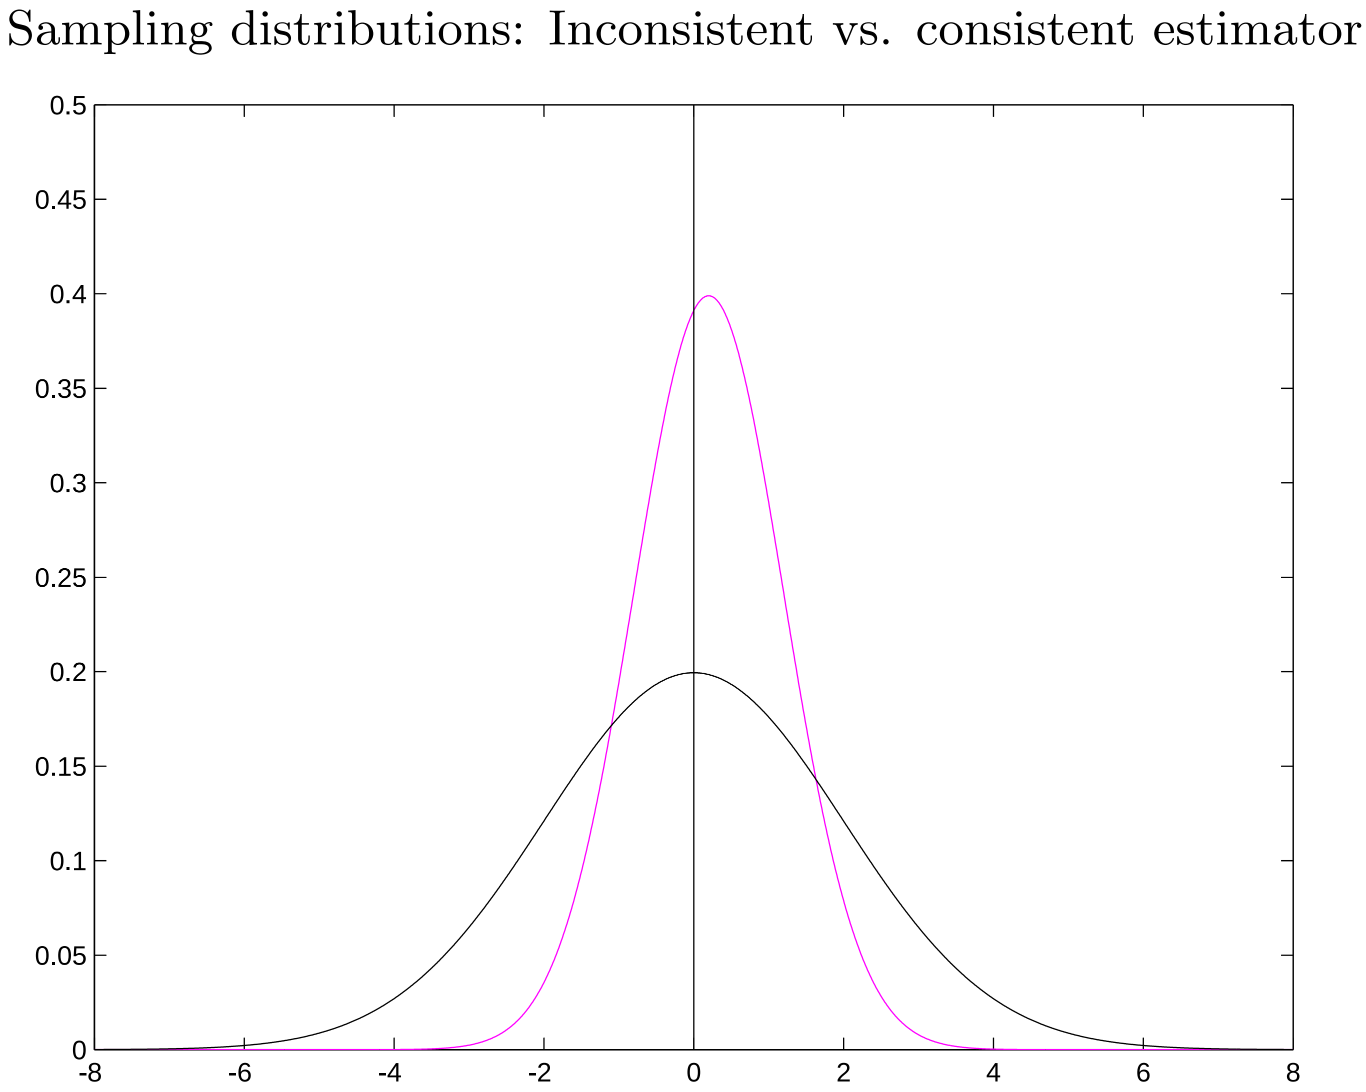

Sampling distributions: Inconsistent vs. consistent estimator
0.5
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
-8
-6
-4
-2
0
2
4
6
8

Motivation

Sampling distributions: Inconsistent vs. consistent estimator (100×T)

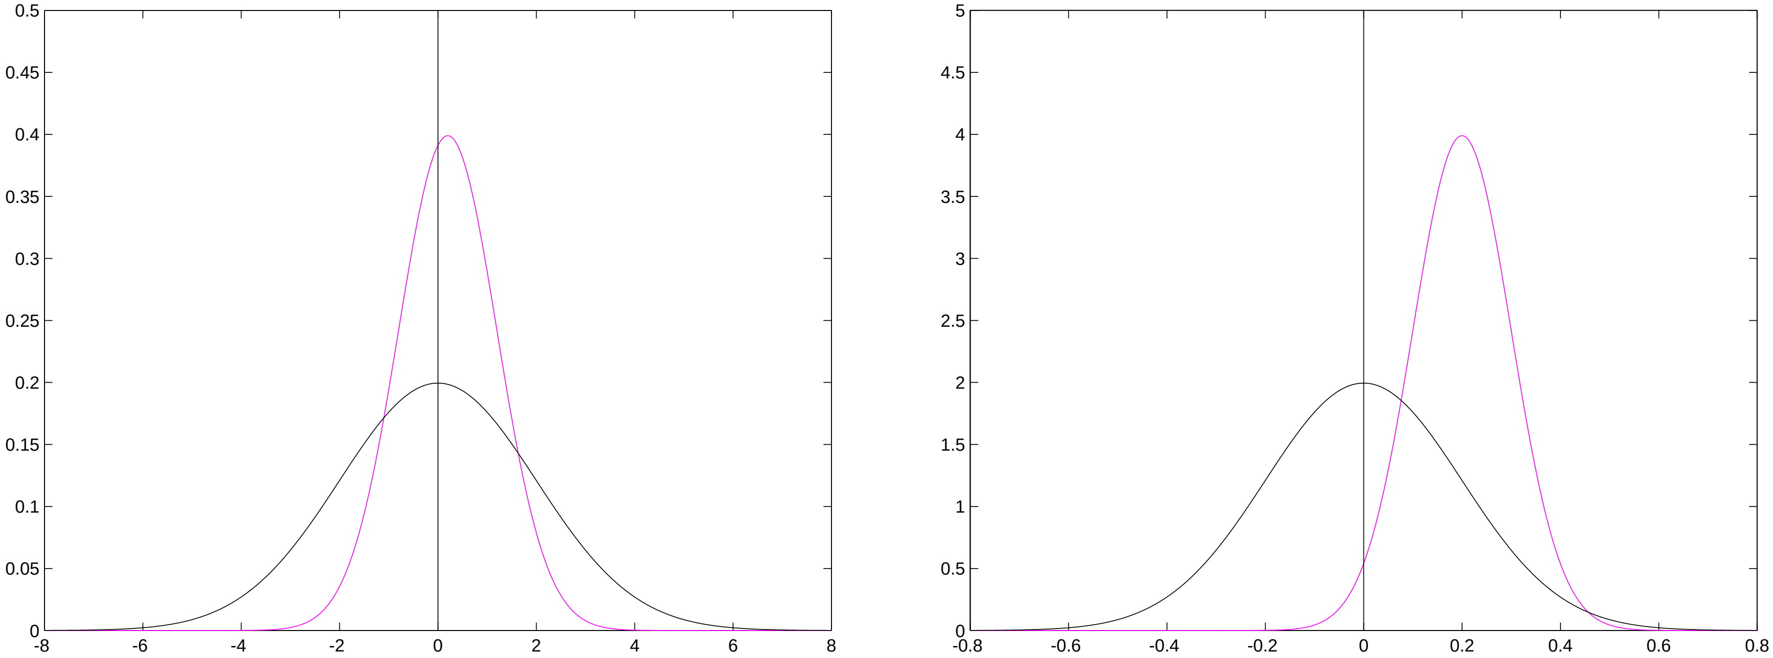

Overview of the talk

1. **Squaring the circle**
2. **Motivation**
3. **Example 1: How profitable is your personal equity plan?**
4. **Example 2: Should you also consider emerging markets?**
5. **Example 3: What happens when things go wrong?**
6. **Conclusions**
7. **Credits**

How profitable is your personal equity plan?

- It's the end of the year, and those of us who're not looking for our first job may wonder how much we should invest in personal equity plans, if at all.
- Those who invested in a plan that tracked the Spanish stock market over the last year should be very happy after seeing their entitlement increased by 30%
- In contrast, those who invested in a guaranteed (i.e. riskless) pension plan may be regretting not having invested in an equity plan.
- But in a standard mean-variance framework it all depends on one's risk preferences and the expected return per unit of risk of equities.

How profitable is your personal equity plan?

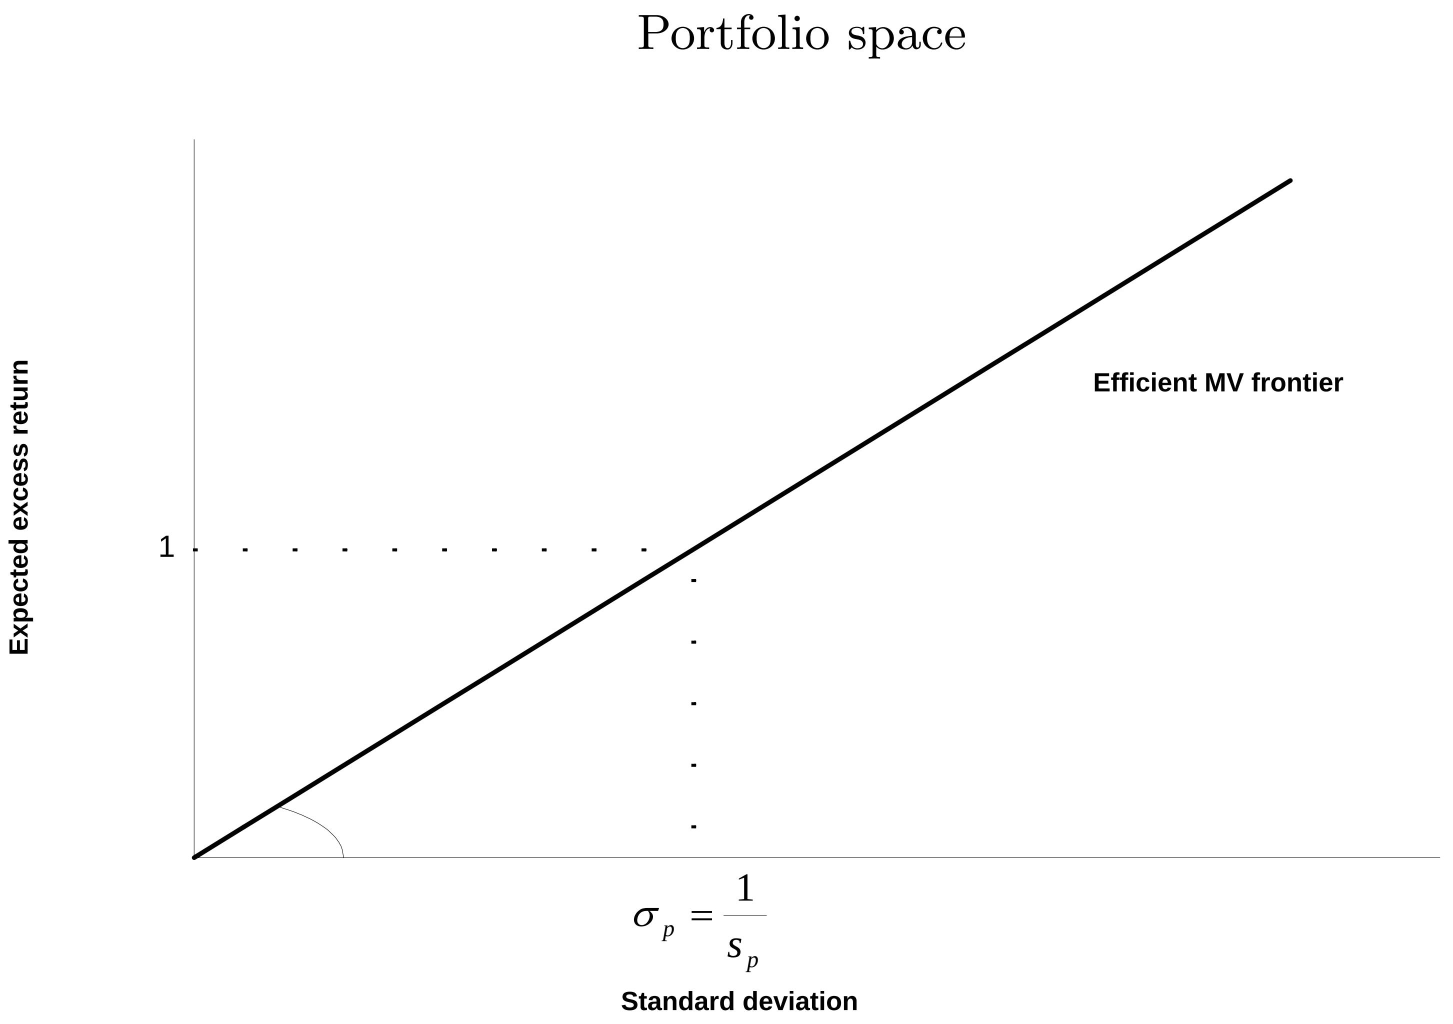

How profitable is your personal equity plan?

Preferences

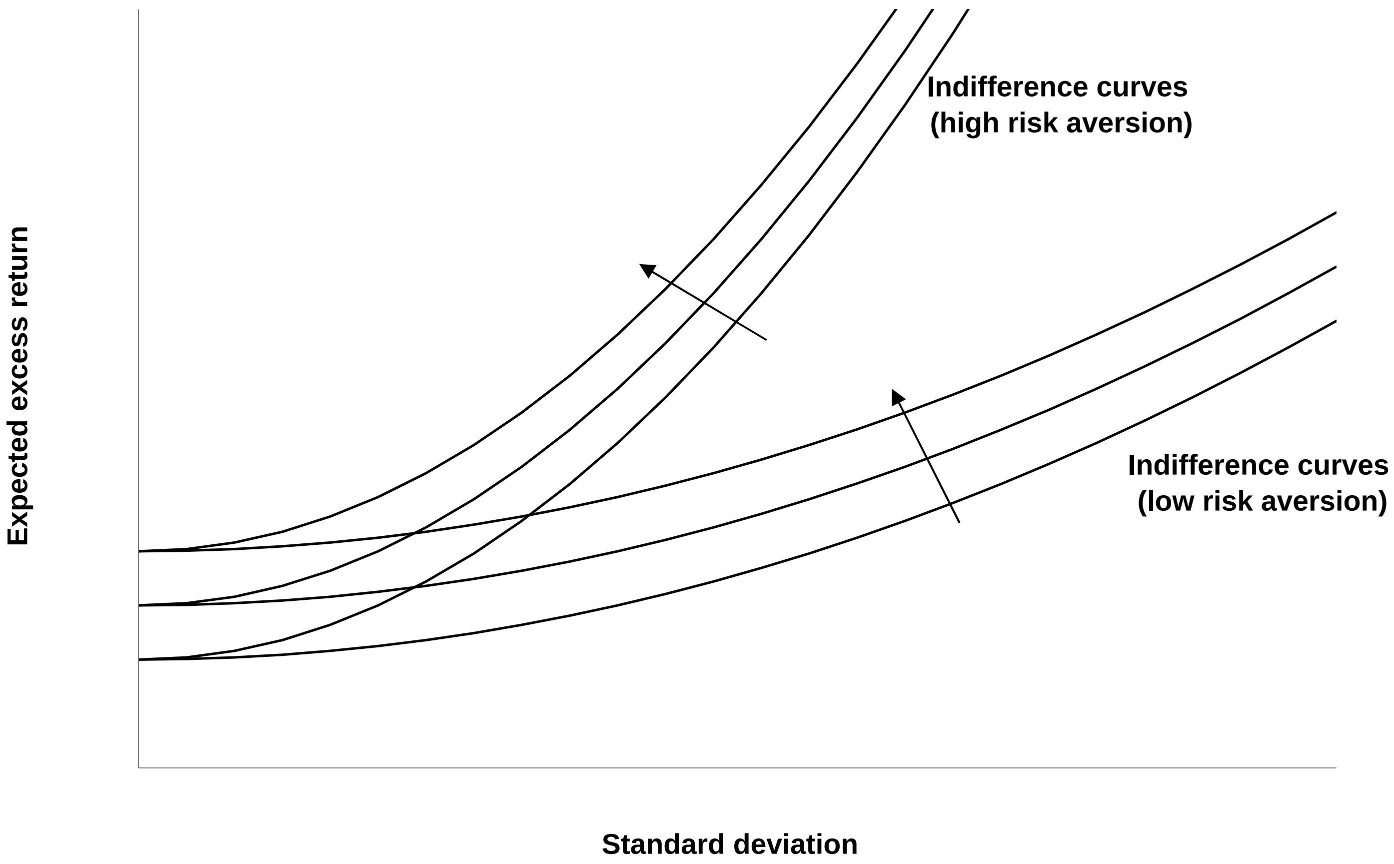

How profitable is your personal equity plan?

Optimal portfolio choice

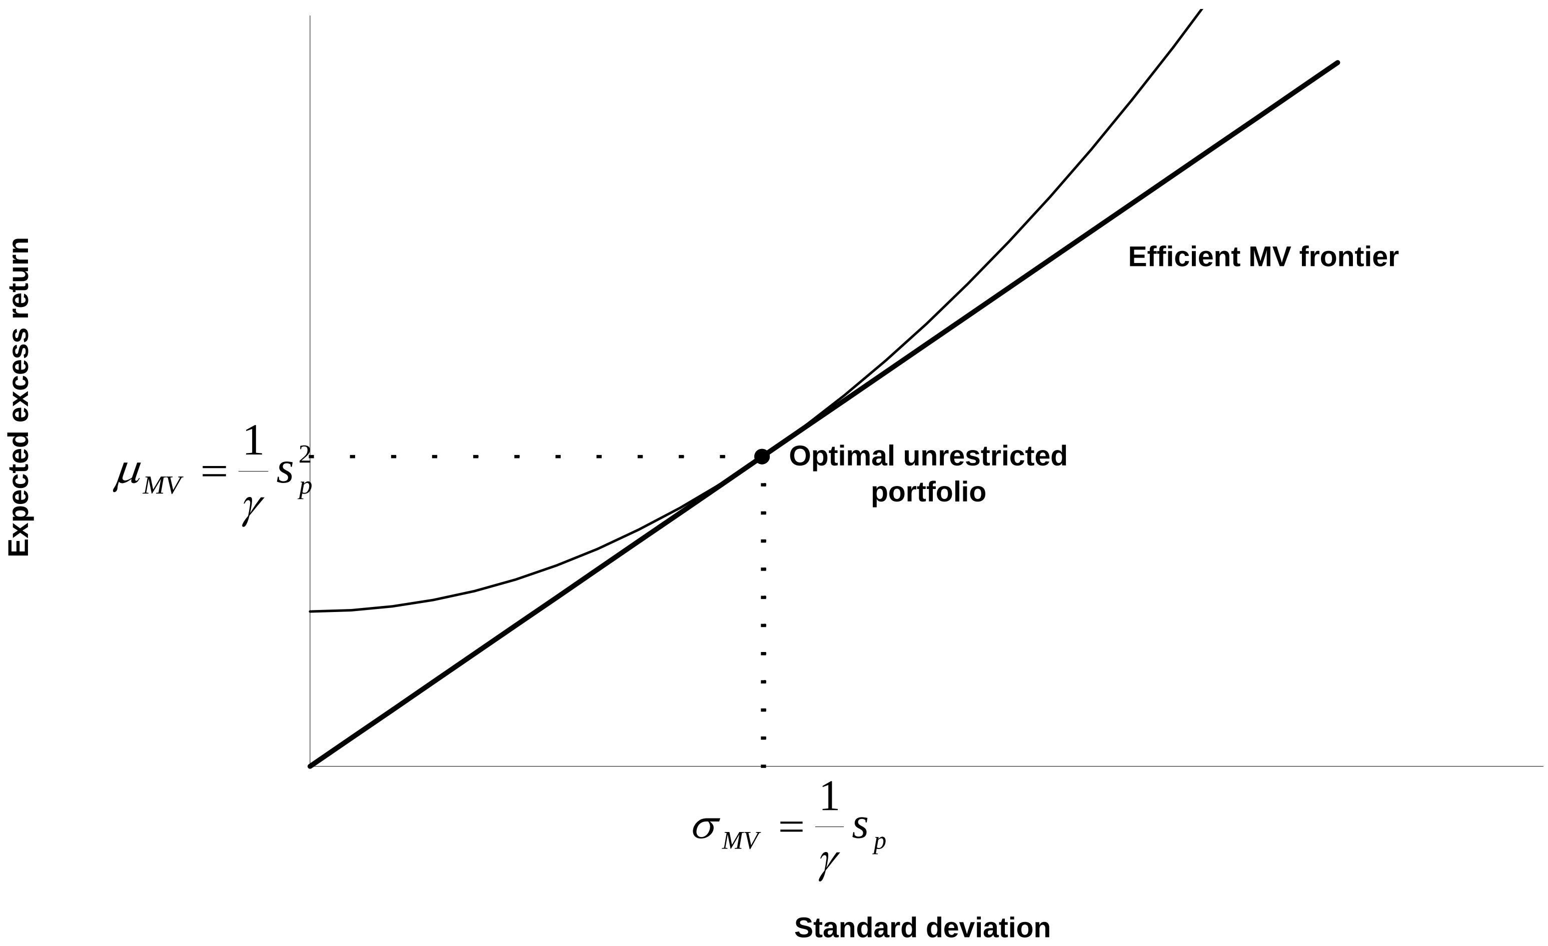

- While we should know our own risk preferences, we don't really know with certainty the true Sharpe ratio of the Spanish stock market $s_m = \mu_m/\sigma_m$.
- Arguably, the most natural estimator is

$$\hat{s}_{mT} = \frac{\hat{\mu}_{mT}}{\hat{\sigma}_{mT}},$$

where:

$$\hat{\mu}_{mT} = \frac{1}{T}\sum_{t=1}^{T} r_{mt},$$

$$\hat{\sigma}^2_{mT} = \frac{1}{T}\sum_{t=1}^{T}(r_{mt} - \hat{\mu}_{mT})^2.$$

- $\hat{s}_{mT}$ is a consistent estimator of $s_m = \mu_m/\sigma_m$, with an asymptotically normal distribution under weak assumptions.

How profitable is your personal equity plan?

- This is not surprising since it is a GMM estimator.
- In fact, it's also both the Gaussian Pseudo-ML estimator and the Semiparametric ML estimator. Therefore, it achieves the semiparametric efficiency bound.
- But unfortunately, it's not very precise.
- To give you an idea of how imprecise it is, I've conducted the following experiment.
- I've assumed that r_{mt} is $i.i.d.$ with annual mean 7% and annual standard deviation 14%, so that the **true** Sharpe ratio is .5 at an annual frequency.
- Then I simulate 10 years of daily data many times, and in each of those parallel universes I compute $\hat{s}_{mT}$, and draw the corresponding mean-variance frontier.

How profitable is your personal equity plan?

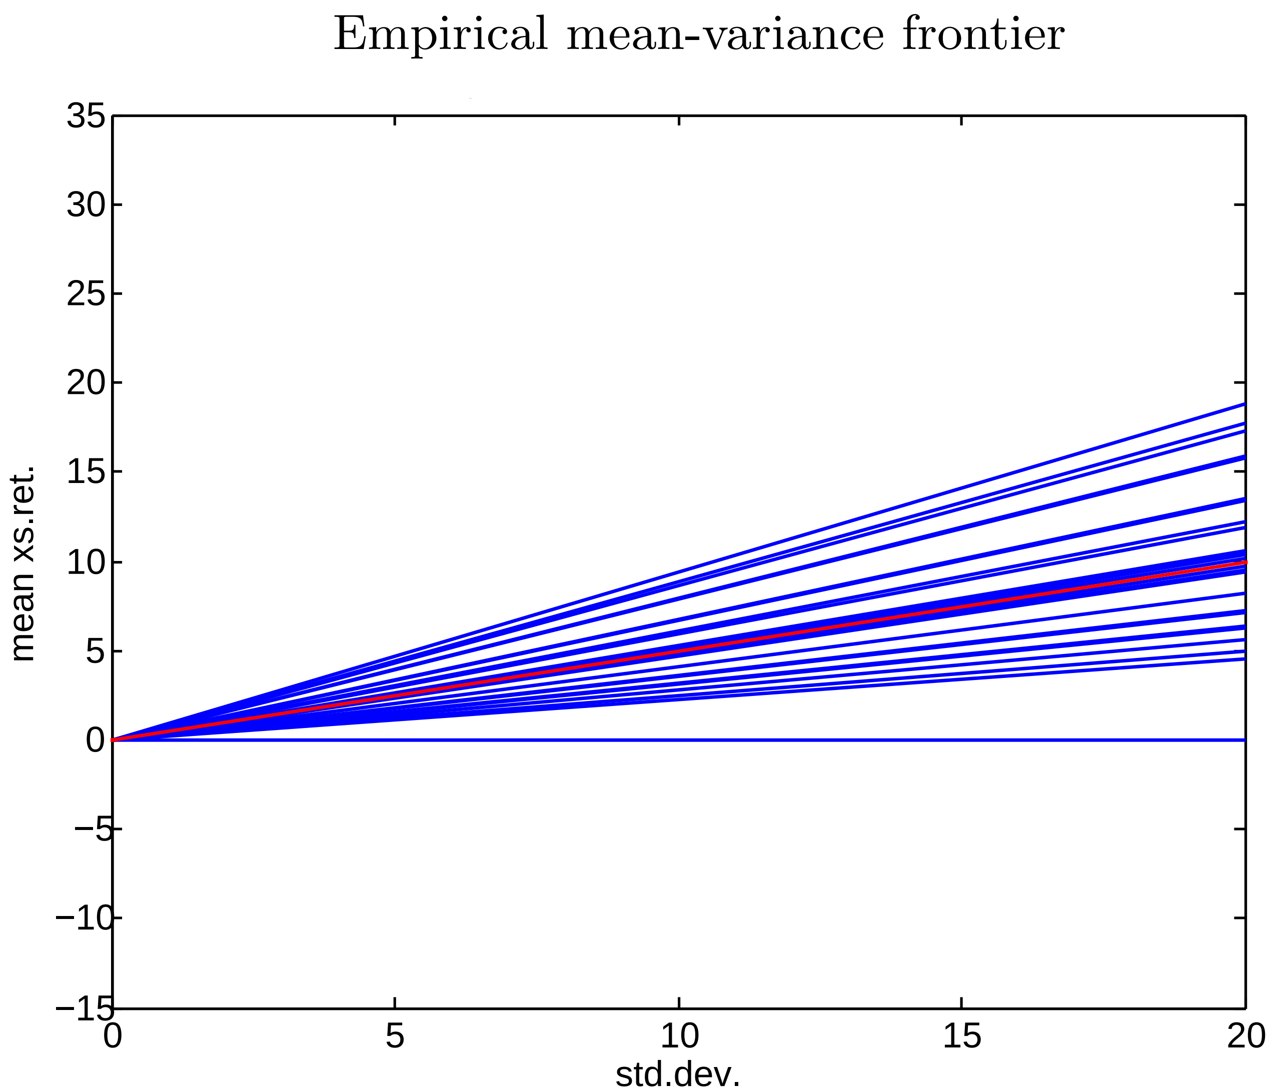

How profitable is your personal equity plan?

Empirical mean-variance frontier

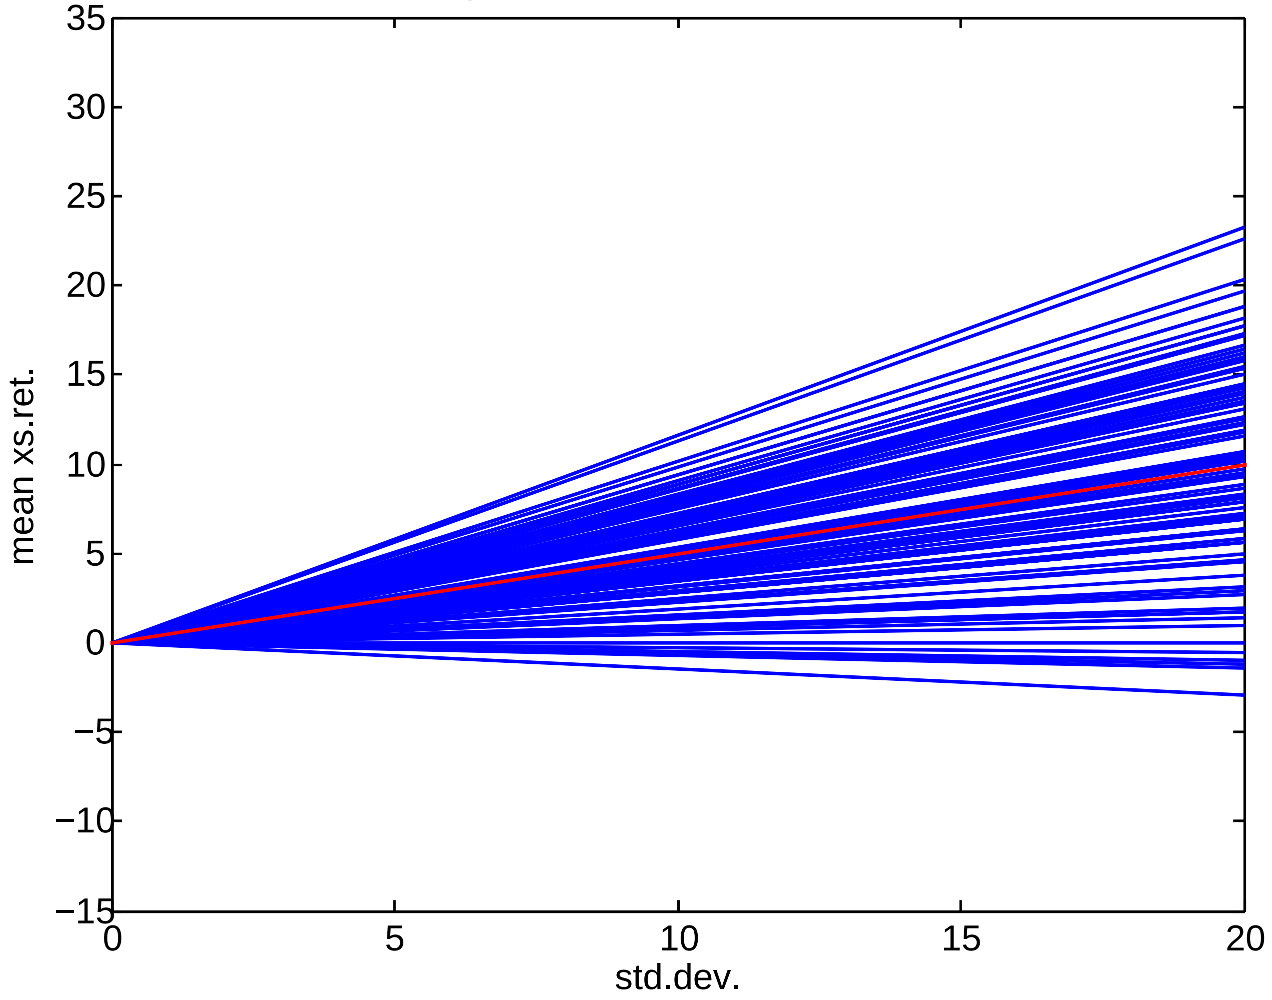

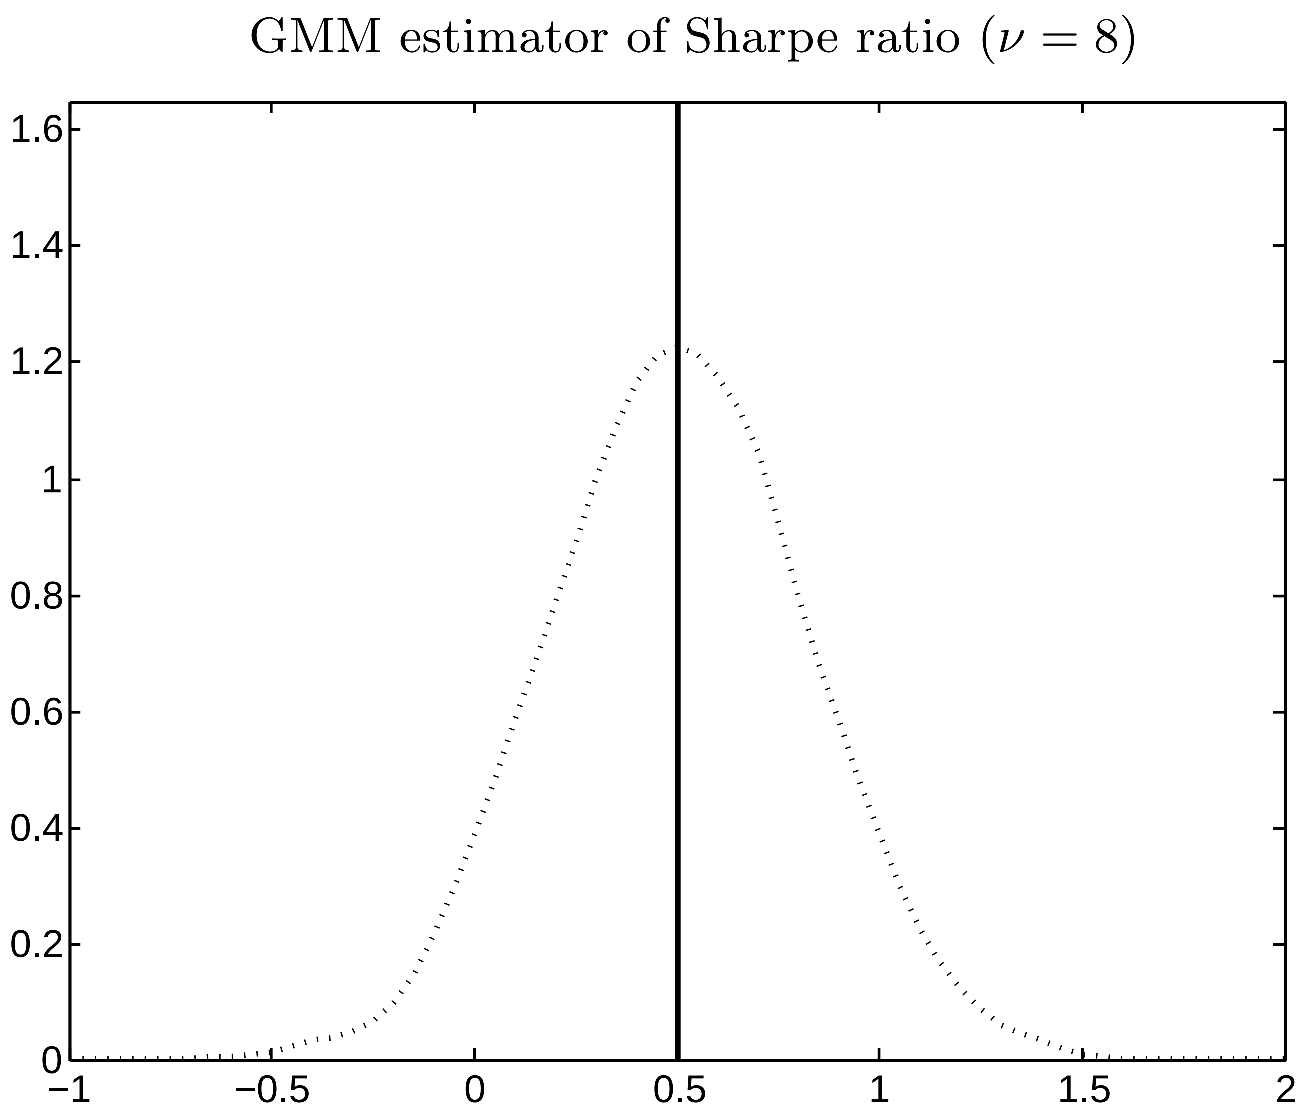
GMM estimator of Sharpe ratio ($\nu = 8$)
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
0
−1
−0.5
0
0.5
1
1.5
2

How profitable is your personal equity plan?

- To reduce sample variability, I've considered two alternative estimators:
  1. Maximum likelihood based on the t distribution
  2. Symmetric semiparametric estimator
- The MLE of μ_m is efficient when the return distribution is t, consistent when it's symmetric, converges to the GMME when it's platykurtic, and is inconsistent otherwise.
- The SSPE of μ_m is efficient when the return distribution is symmetric, and inconsistent otherwise.
- The MLE of σ_m^2 is efficient when the return distribution is t, converges to the GMME when it's platykurtic, and is inconsistent otherwise, although it can be bias-corrected under symmetry.
- The SSPE of σ_m^2 converges to the GMME when the return distribution is symmetric, and is inconsistent otherwise.

How profitable is your personal equity plan?

Sharpe ratio estimators ($\nu = 8$)

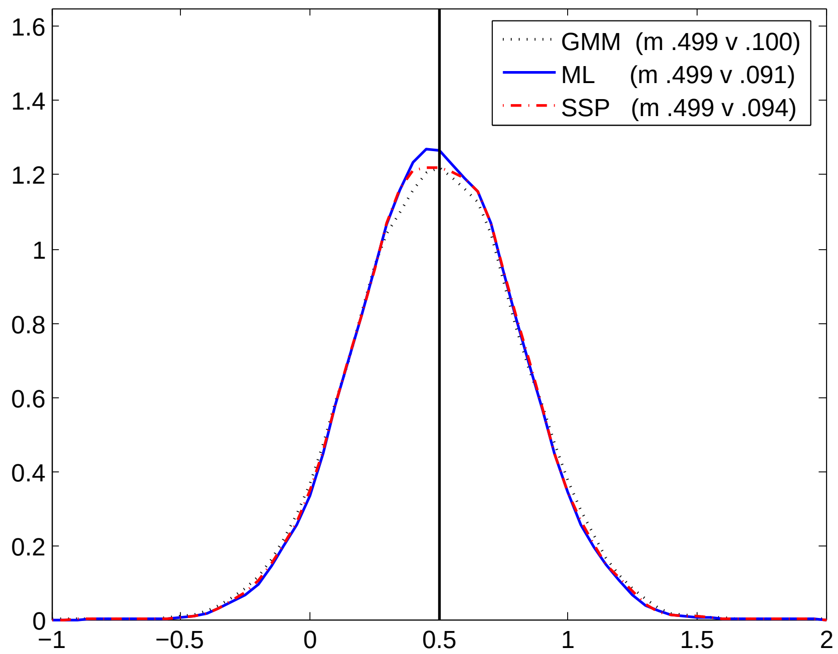

Student t ($\nu = 8$)

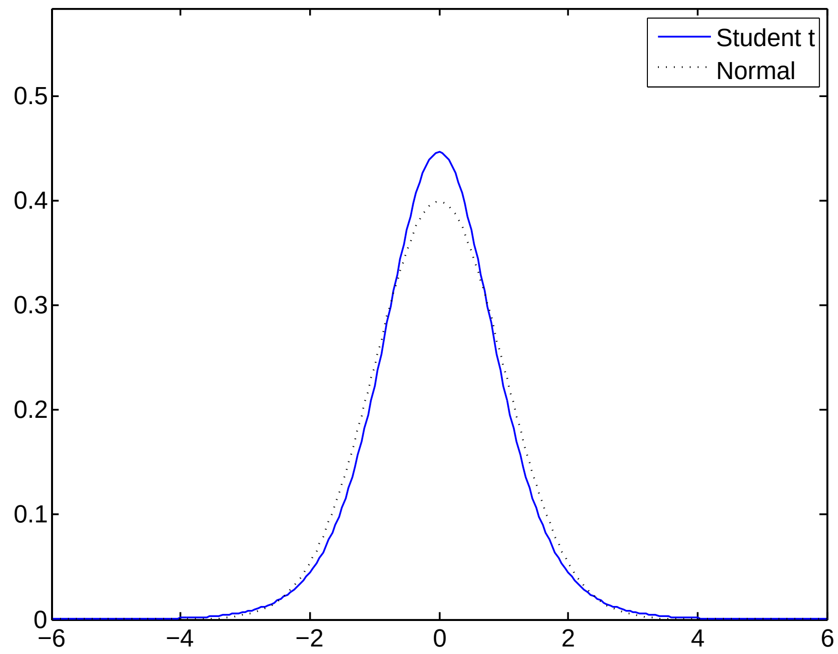

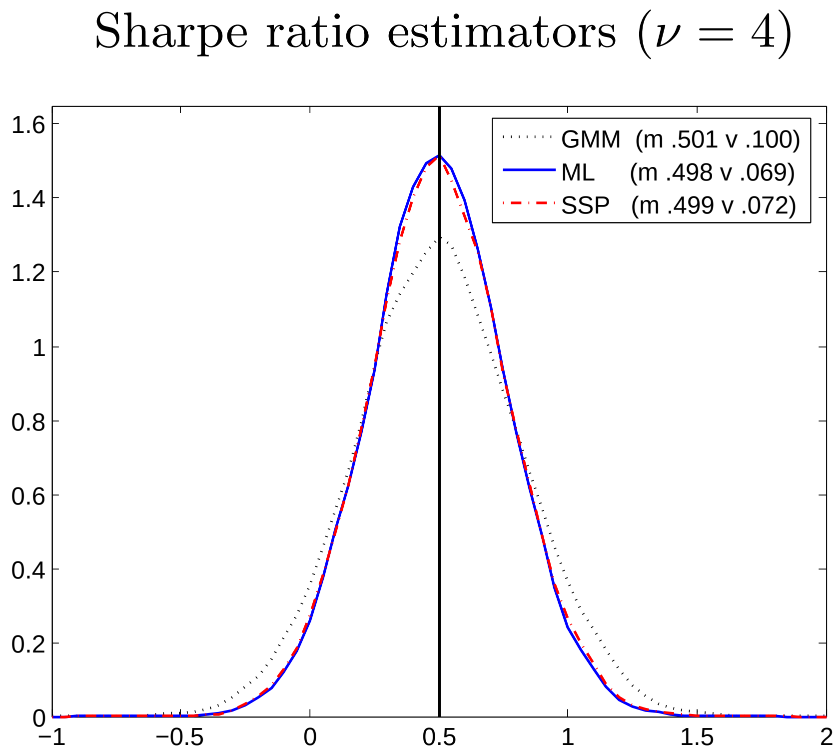
Sharpe ratio estimators ($\nu = 4$)
GMM (m .501 v .100)
ML (m .498 v .069)
SSP (m .499 v .072)
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
0
−1
−0.5
0
0.5
1
1.5
2

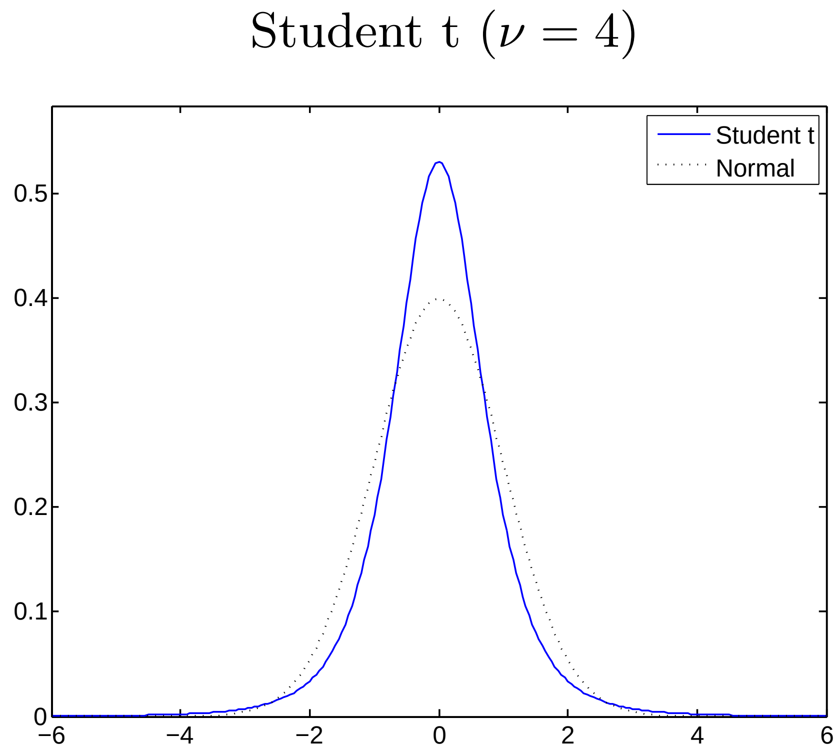
Student t ($\nu = 4$)
Student t
Normal
0.5
0.4
0.3
0.2
0.1
0
−6
−4
−2
0
2
4
6

How profitable is your personal equity plan?

Sharpe ratio estimators (normal-gamma)

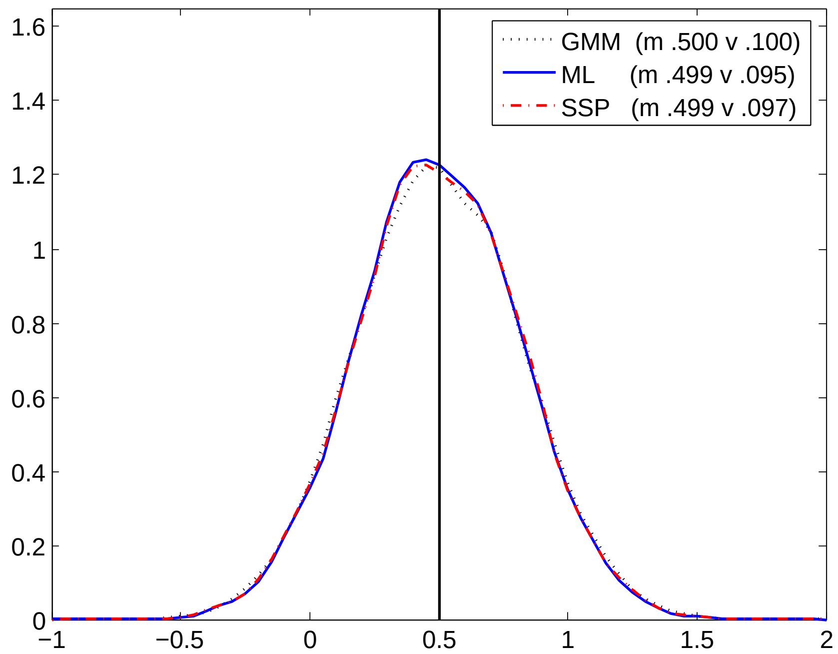

Symmetric normal-gamma mixture

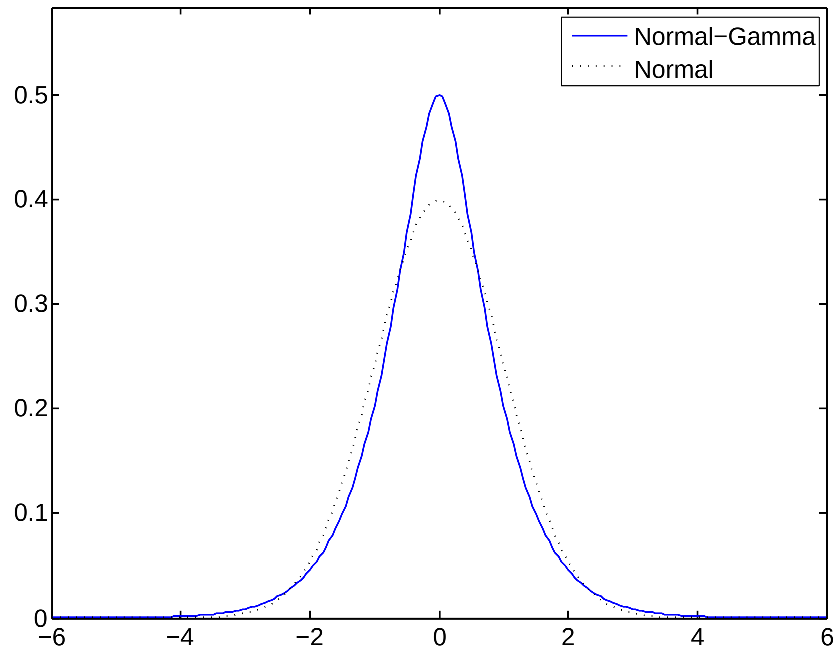

How profitable is your personal equity plan?

Sharpe ratio estimators (asymmetric t)

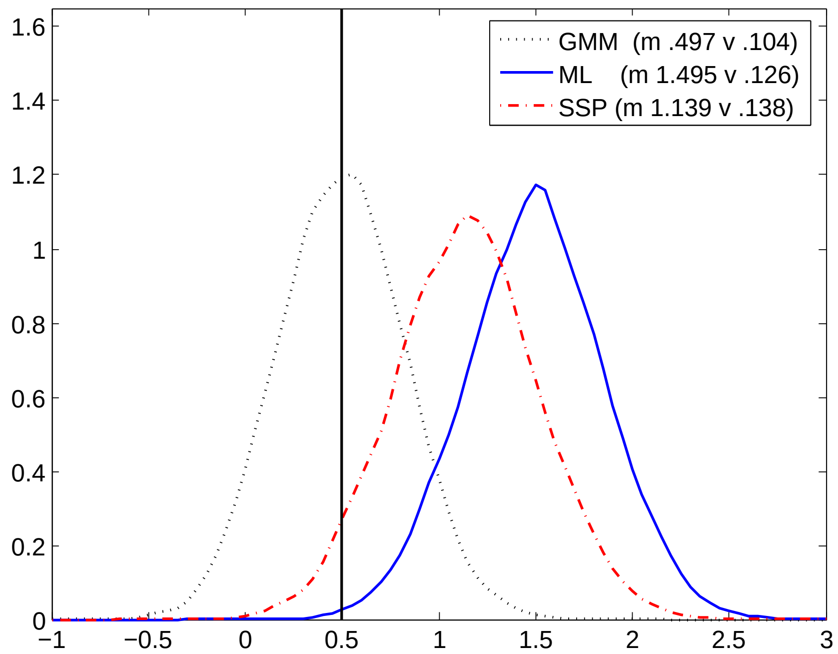

Asymmetric t

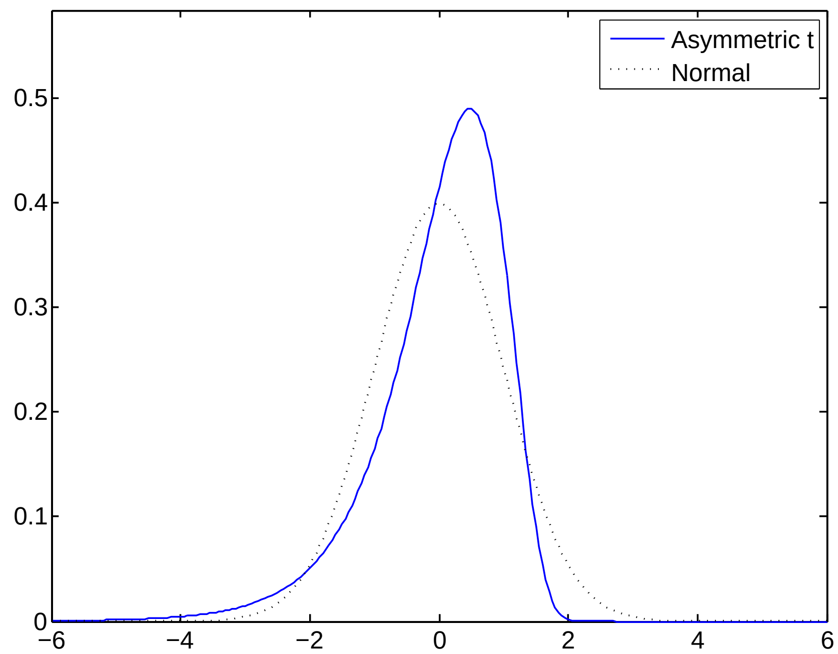

Should you also invest in emerging markets?

- You've now decided to invest in a personal equity plan, but when you go to your bank, your investment advisor, noticing your good disposition, tries to convince you to invest in BRIC equities. (Note: BRIC $\equiv$ Brasil, Russia, India, China)
- She tells you that despite their higher volatility, they are sound investments because:
  1. they offer a commensurate expected rate of return,
  2. they provide huge diversification benefits because they are uncorrelated with the Spanish stock market.
- But is it really true?
- You go home and download data from the internet. Unfortunately, you can only get the last two years of daily data.
- What can you do to decide?

Should you also invest in emerging markets?

- The first thing to note is that you shouldn't simply look at whether the slope of the empirical mean-variance frontier increases, because it always does, even if the theoretical frontier is unaffected.
- To convince you that what I'm saying is true I've conducted the following experiment:
- I've assumed that you can invest in r_{mt} and in three additional assets, r_{1t}, r_{2t}, and r_{3t}, which are *i.i.d.* with an annual mean of 0% and annual standard deviation of 28%, uncorrelated among themselves and with the original asset, so that the **true** Sharpe ratio doesn't increase.
- Then I simulate 2 years of daily data many times, and compute the original and augmented mean-variance frontiers, as well as the incremental one, which corresponds to the differences between r_{1t}, r_{2t}, and r_{3t} and their best tracking portfolios based on r_{mt}.

Should you also invest in emerging markets?

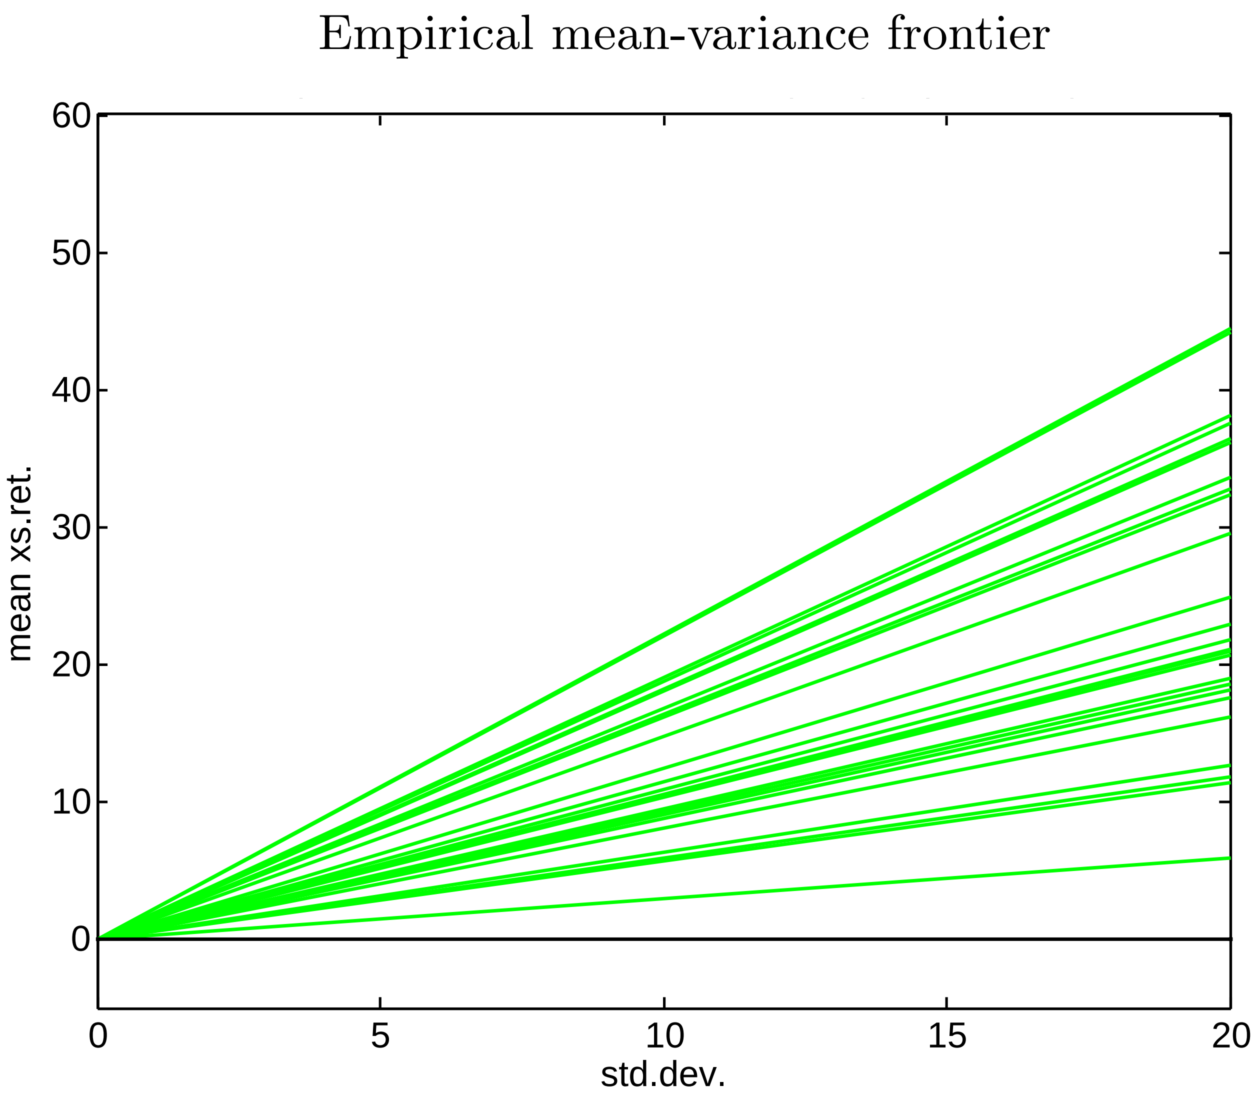

Should you also invest in emerging markets?

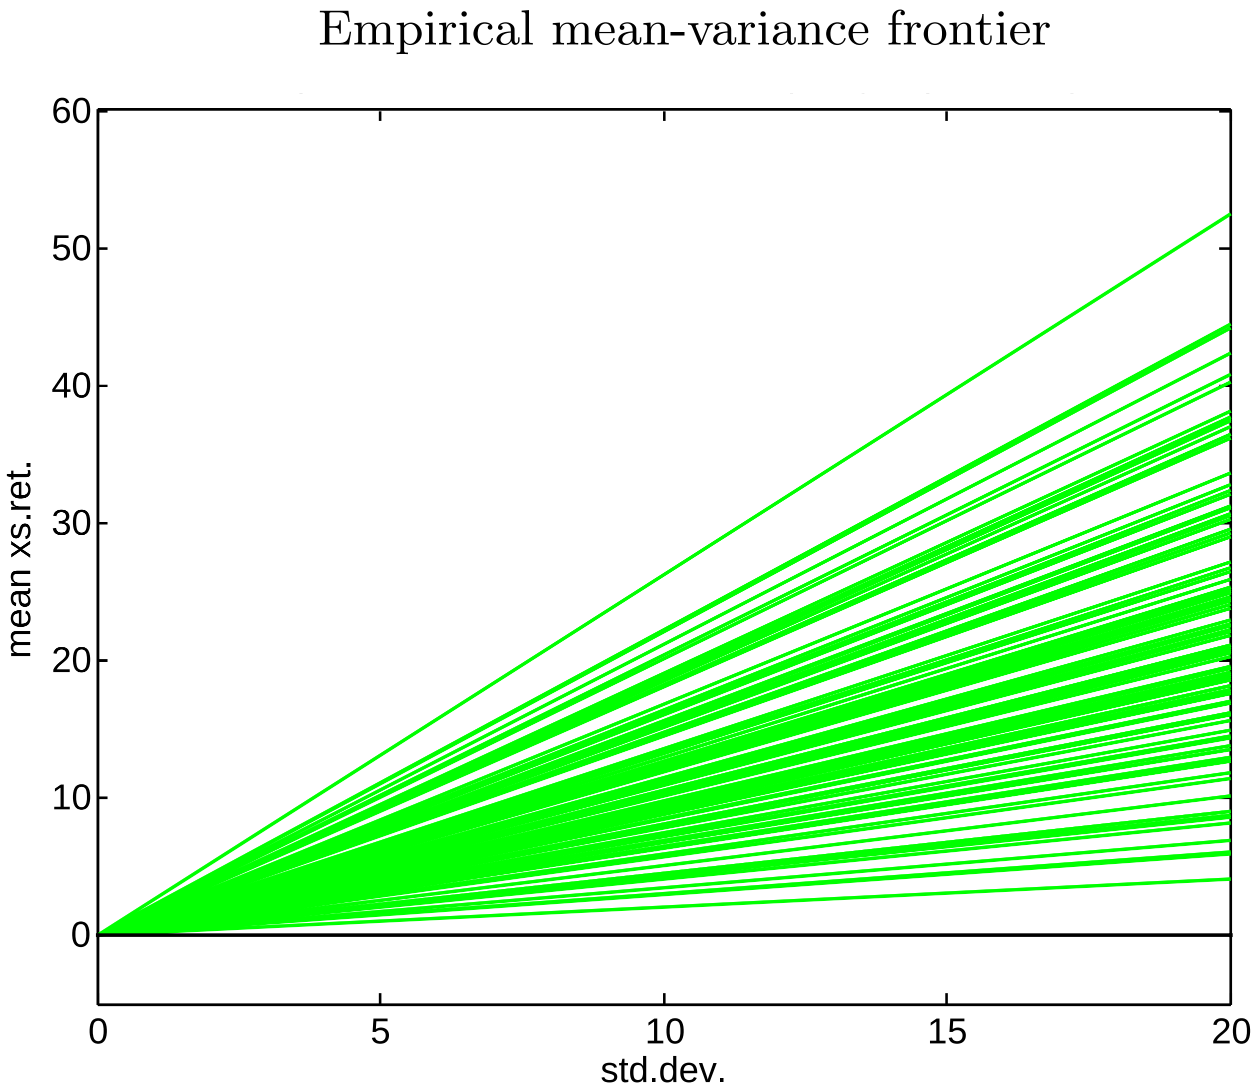

Should you also invest in emerging markets?

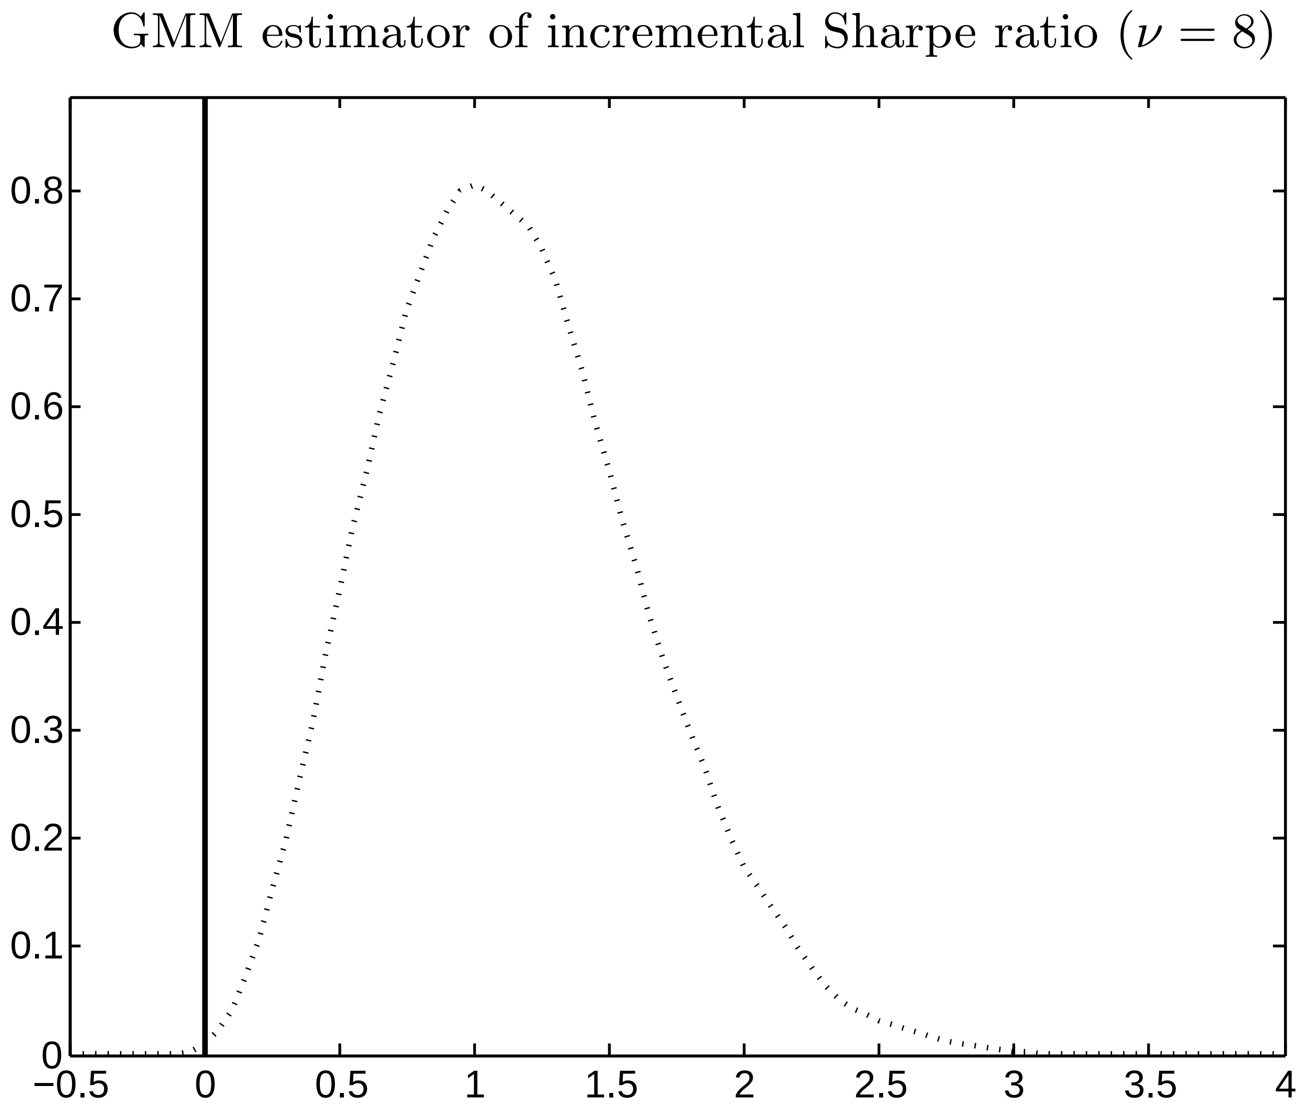

Should you also invest in emerging markets?

- Instead, you should look at the means of the hedged portfolios

$$\alpha_i = \mu_i - \frac{\sigma_{im}}{\sigma_m^2}\mu_m = \mu_i - \beta_{im}\mu_m,$$

where

$$\mu_i = E(r_{it}),$$

$$\sigma_{im} = cov(r_{it}, r_{mt}),$$

and test whether $H_0 : \alpha_i = 0$, $(i = 1, 2, 3)$.

- You could do this with GMM by writing down the multivariate version of the normal equations (f.o.c.) as moment conditions:

$$E\left\{\begin{pmatrix} 1 \\ r_{mt} \end{pmatrix} \otimes \left[\begin{pmatrix} r_{1t} \\ r_{2t} \\ r_{3t} \end{pmatrix} - \begin{pmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \end{pmatrix} - \begin{pmatrix} \beta_{1m} \\ \beta_{2m} \\ \beta_{3m} \end{pmatrix} r_{mt}\right]\right\} = 0.$$

Should you also invest in emerging markets?

- But there are other, potentially more efficient, alternatives:
  1. Maximum likelihood tests based on the assumption that the distribution of r_{1t}, r_{2t} and r_{3t} conditional on r_{mt} is a multivariate student t, which can be easily robustified when the true distribution is elliptical.
  2. Symmetric semiparametric tests based on the assumption that this conditional distribution is elliptical.
- Elliptical distributions generalise the multivariate normal distribution, but at the same time they retain its analytical tractability irrespective of the number of assets.
- Moreover, they guarantee that mean-variance analysis is fully compatible with expected utility maximisation regardless of investors' preferences.

Should you also invest in emerging markets?

Elliptical densities

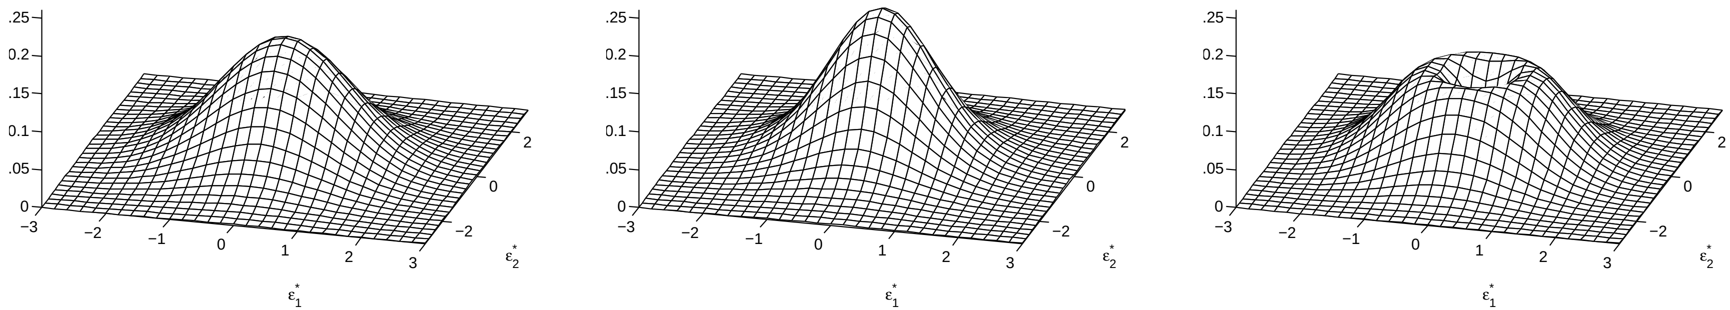

Contours

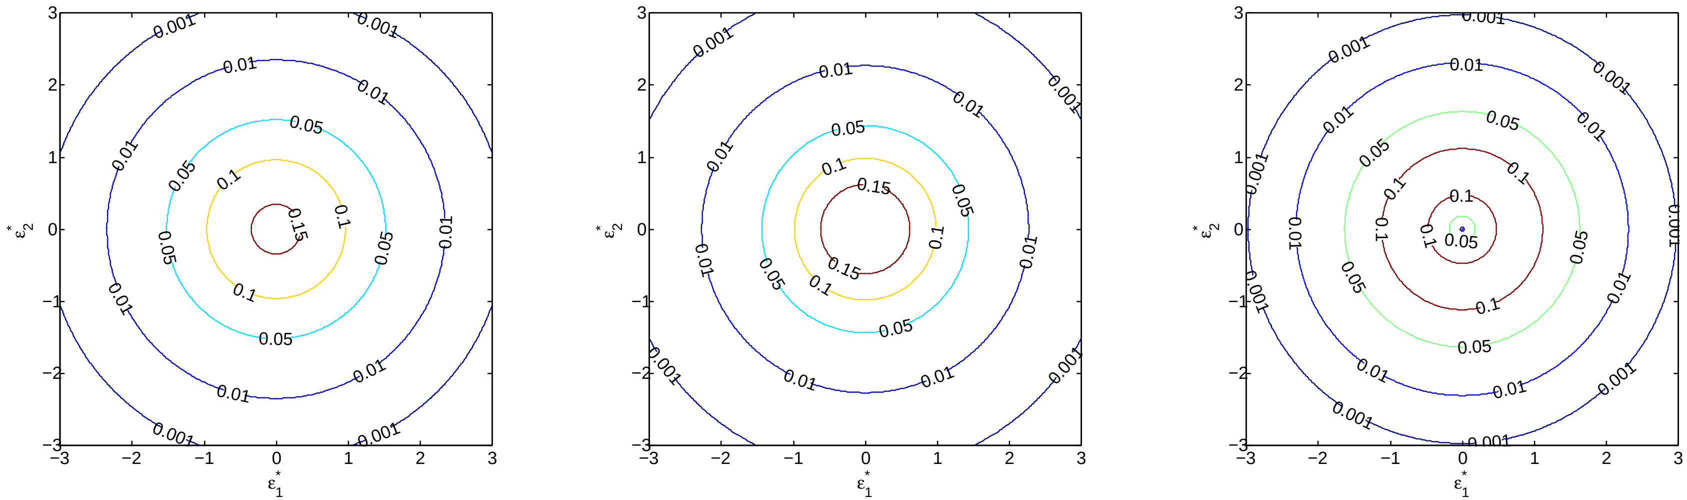

Normal | Student t | Kotz

Trivariate student t $\quad (\nu = 8)$

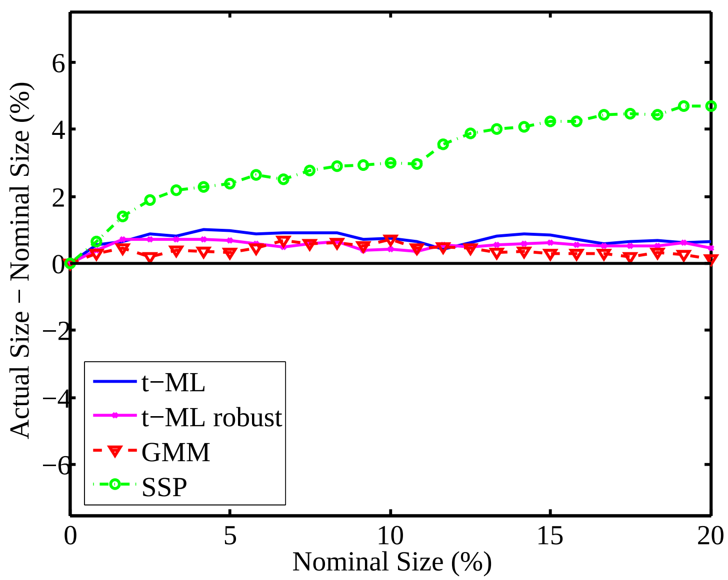

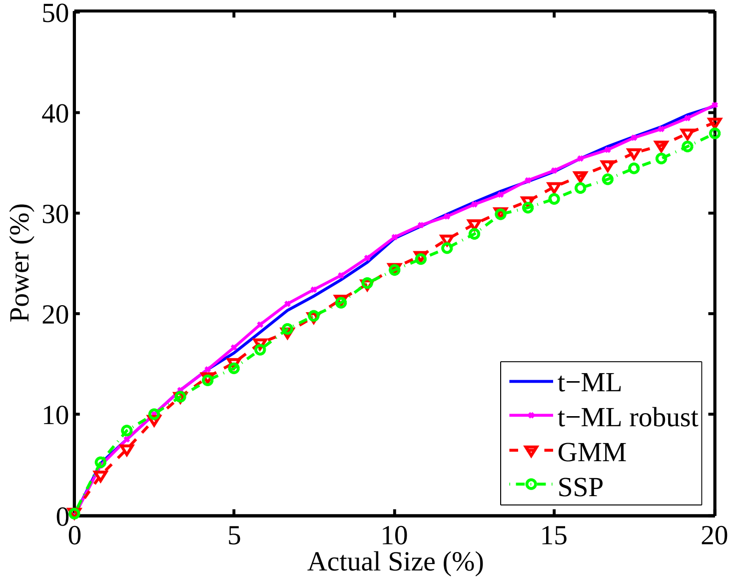

Trivariate student t $\quad (\nu = 4)$

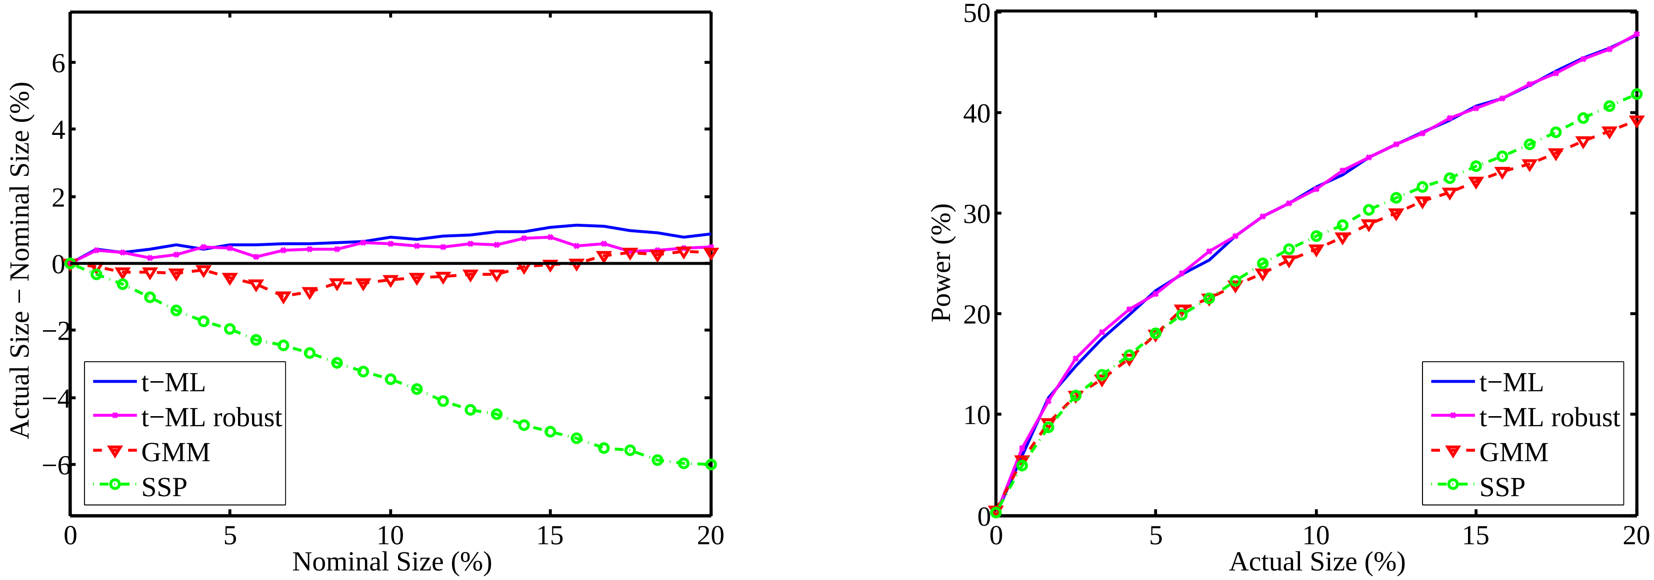

Trivariate normal inverse Gaussian

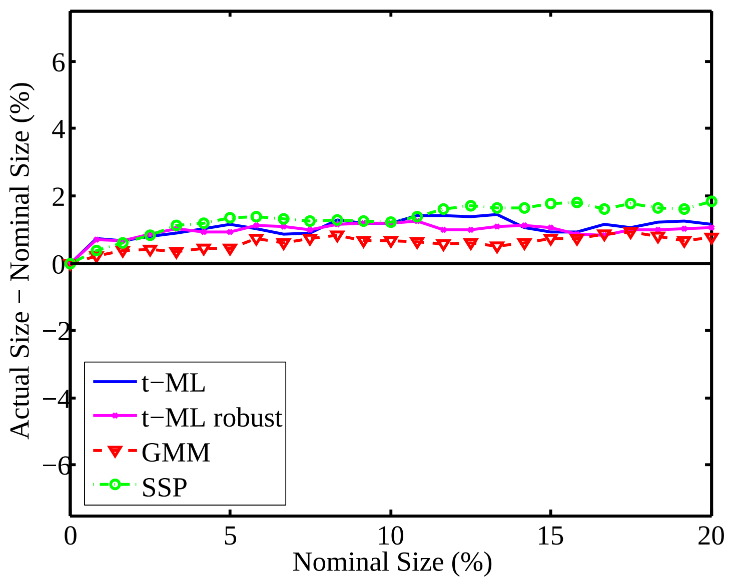

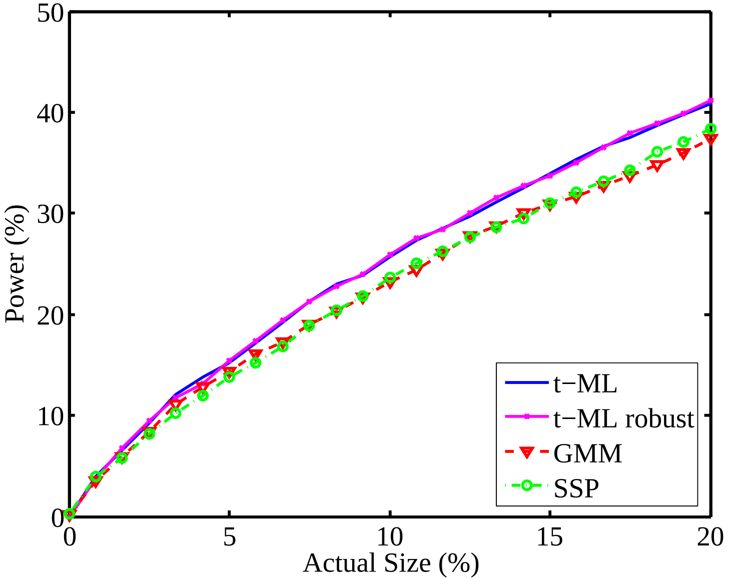

Should you also invest in emerging markets?

Trivariate asymmetric t

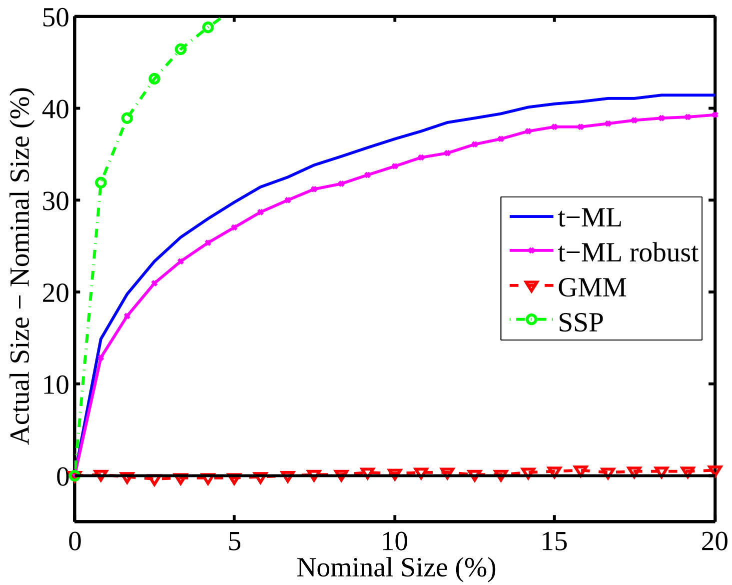

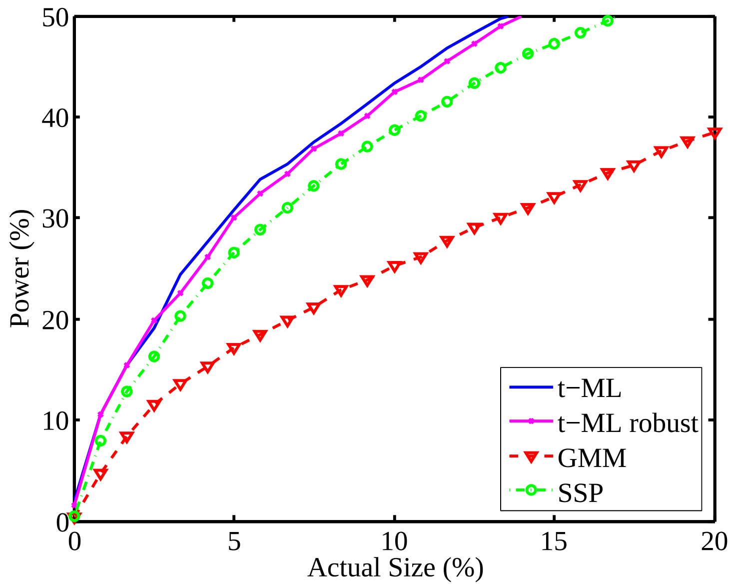

What happens when things go wrong?

- Imagine that you've now decided to invest in one of the BRIC countries.
- Unfortunately, financial markets that show little correlation in normal periods have a tendency to fall together during crisis periods.
- For instance, when Russia defaulted in August 1998, most stock markets around the world fell at unison.
- This is really terrible if you thought that you could eliminate most of the risk of your portfolio by holding assets that are seemingly independent (i.e. uncorrelated).
- Of course, there are several ways in which one could model **'tail dependence'**.
- But in any case, such a phenomenon is something that the multivariate normal distribution cannot account for, and elliptical distributions struggle with.

What happens when things go wrong?

The Generalised Hyperbolic distribution is a multivariate infinitely divisible distribution that nests as particular cases:

- Gaussian,
- Student t,
- Asymmetric Student t,
- Normal-Gamma mixture,
- Normal Inverse Gaussian,

and several other popular distributions.

What happens when things go wrong?

It can be generated as the following location-scale mixture of normals

$$\boldsymbol{\varepsilon}^{*}=\boldsymbol{\alpha}+\boldsymbol{\Upsilon}\boldsymbol{\beta}\xi^{-1}+\xi^{-\frac{1}{2}}\boldsymbol{\Upsilon}^{\frac{1}{2}}\mathbf{z},$$

where

$$\mathbf{z}\sim N(\mathbf{0},\mathbf{I}_N),$$

$$\xi\sim GIG\,(-\nu,\gamma,\delta),$$

so that

$$\boldsymbol{\varepsilon}^{*}|\xi\sim N(\boldsymbol{\alpha}+\boldsymbol{\Upsilon}\boldsymbol{\beta}\xi^{-1},\boldsymbol{\Upsilon}\xi^{-1}).$$

- $\boldsymbol{\alpha}$: location parameter,
- $\boldsymbol{\Upsilon}$: dispersion matrix,
- δ: scale parameter,
- ν and γ: tail shape parameters,
- $\boldsymbol{\beta}$: skewness parameter.

What happens when things go wrong?

- I can always choose $\boldsymbol{\alpha}$ and $\boldsymbol{\Upsilon}$ so that the elements of $\boldsymbol{\varepsilon}^*$ are uncorrelated but not independent, with zero means and unit variances.
- In general, the GH distribution induces "tail dependence" through ξ, which can be asymmetrically calibrated with $\boldsymbol{\beta}$.
- In this way, we can capture the empirically observed higher tail dependence across stock returns in market downturns.

What happens when things go wrong?

GH densities

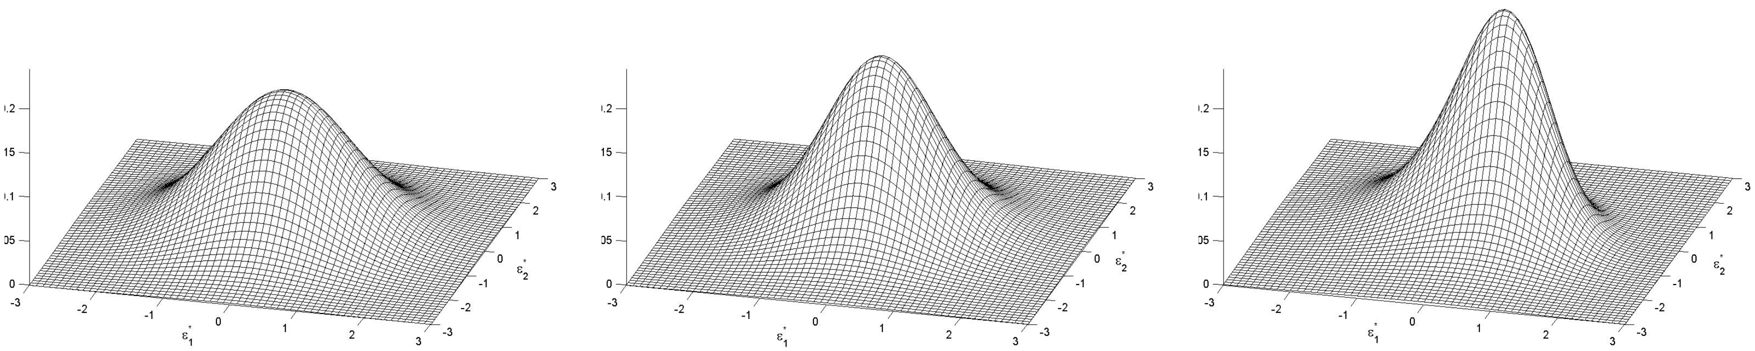

Contours

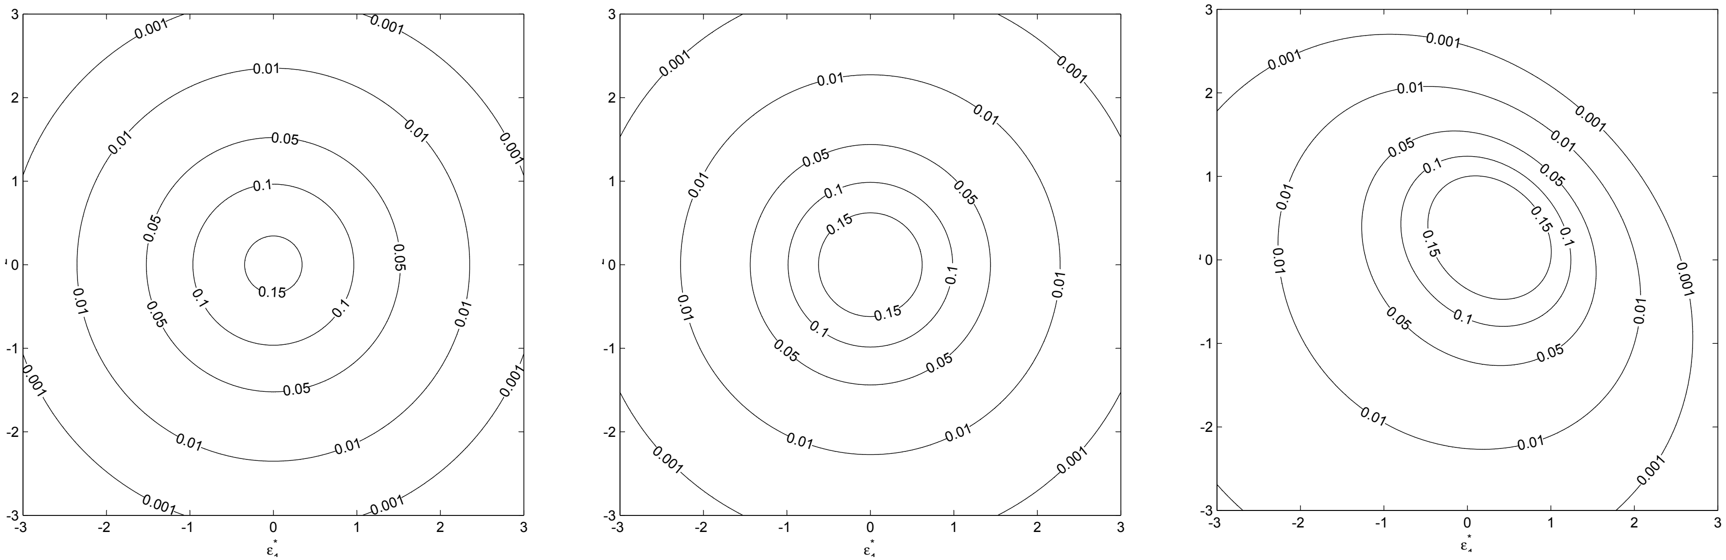

Normal   Student t Asymmetric t

$\eta = .1$ $\boldsymbol{\beta} = (-3, -3)'$

What happens when things go wrong?

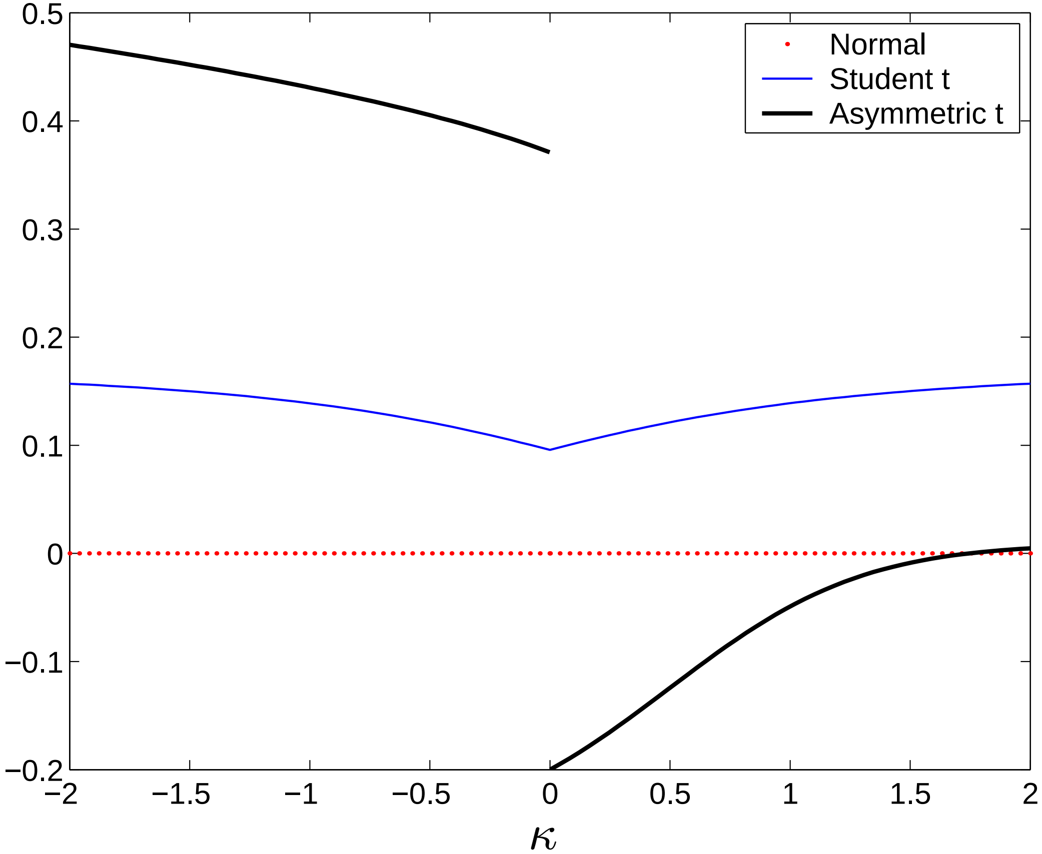

$$\text{Definition: } \bar{\rho}(\kappa) = \begin{cases} \operatorname{corr}(\varepsilon_1^*, \varepsilon_2^* | \, \varepsilon_1^* > \kappa, \varepsilon_2^* > \kappa) & \text{if} \quad \kappa \geq 0 \\ \operatorname{corr}(\varepsilon_1^*, \varepsilon_2^* | \, \varepsilon_1^* < \kappa, \varepsilon_2^* < \kappa) & \text{if} \quad \kappa < 0 \end{cases}$$

What happens when things go wrong?

- In principle, we could be fully semi-parametric, and estimate

$$\mathrm{corr}(\varepsilon_1^*, \varepsilon_2^* | \varepsilon_1^* < \kappa, \varepsilon_2^* < \kappa)$$

  by computing the sampling analogue to this definition after having estimated the means and standard deviations of the observed variables.

- The problem is that for $|\kappa|$ relatively large, say 2.5 standard deviations, there are very few observations in the appropriate region, as severe crises are rare (.37% in the case of the asymmetric t, .0038% for the normal).

- In those circumstances, a flexible distribution such as the GH offers a very useful alternative.

Conclusions

- Semiparametric methods such as GMM, together with fully nonparametric procedures, have revolutionised the practice of Econometrics.
- However, standard (i.e. $T^{\frac{1}{2}}$) asymptotic theory is often not very reliable as an approximation to their finite sample behaviour.
- Maximum likelihood procedures based on a flexible parametric distribution can provide useful alternatives if judiciously chosen.
- We need an alternative limiting theory to analyse the consistency-efficiency trade-offs that they offer.
- In addition, there are situations in which we are interested in characteristics of a distribution that go beyond its moments.
- In those cases, a flexible parametric distribution can also prove very useful.

Credits

1. **How profitable is your personal equity plan?**
   Fiorentini and Sentana (2006): 'The relative efficiency of pseudo-maximum likelihood estimation and inference in conditionally heteroskedastic dynamic regression models'

2. **Should you also consider emerging markets?**
   Amengual and Sentana (2006): 'Asymptotic comparisons of mean-variance efficiency tests'

3. **What happens when things go wrong?**
   Mencía and Sentana (2005): 'Estimation and testing of dynamic models with generalised hyperbolic innovations'