Parameter Uncertainty and Solvency Risk

Andreas Tsanakas

Cass Business School, City University London

based on joint work with R. Gerrard, V. Bignozzi, Z. Landsman

Actuarial & Financial Mathematics 2012: Theory & Applications

University of Liverpool, 7 June 2012

Motivation

- Regulation
  - Strict requirements (eg 99.5%-VaR)
  - Limited data...
- What is this talk about?
  - An idiosyncratic tour of questions that have troubled me
  - Not a comprehensive review
  - Omitting much mathematical detail

Setup (1)

- I.i.d. data $X_1, \ldots, X_n \sim F(\cdot; \theta)$
- Future loss $Y \sim F(\cdot; \theta)$, independent of $\mathbf{X} = \{X_1, \ldots, X_n\}$
- Parameter(s) θ unknown, to be estimated by MLE $\widehat{\theta} = \widehat{\theta}(\mathbf{X})$
- MLE of distribution: $F(\cdot; \widehat{\theta})$
- Predictive distribution: $\widehat{F}(\cdot|\mathbf{X}) = \int F(\cdot; t)\pi(t|\mathbf{X})dt$

Setup (2)

- Regulatory risk measure ρ
  - Monotonic: $Y \leq Z \implies \rho(Y) \leq \rho(Z)$
  - Translation invariant: $\rho(Y - \lambda) = \rho(Y) - \lambda$
  - Positive homogenous: For $\lambda > 0$, $\rho(\lambda Y) = \lambda \rho(Y)$
  - Law invariant: If $Y \stackrel{d}{=} Z$, then $\rho(Y) = \rho(Z)$.
    - So for $Y \sim F(\cdot;\theta)$, let $\rho_\theta(Y) = \rho[F(\cdot;\theta)]$

- For example

$$VaR_{p,\theta}(Y) = VaR_p[F(\cdot;\theta)] = \inf\{y \in \mathbb{R} : F(y;\theta) \geq p\}$$

$$TVaR_{p,\theta}(Y) = TVaR_p[F(\cdot;\theta)] = \frac{1}{1-p}\int_p^1 VaR_{\alpha,\theta}(Y)d\alpha$$

Parameter uncertainty

- $\rho_\theta(Y)$ is the minimum level of capital such that

  $\rho_\theta\left(Y - \rho_{\theta}(Y)\right) = 0$

- In reality capital held is $\rho_{\widehat{\theta}}(Y) = \rho[F(\cdot;\widehat{\theta})]$, such that

  $\rho_\theta\left(Y - \rho_{\widehat{\theta}}(Y)\right) \neq 0$

- Randomness of capital reflects variation across time periods or firms

- When $\rho_\theta\left(Y - \rho_{\widehat{\theta}}(Y)\right) \geq 0$ a positive *residual risk* is incurred (Bignozzi & T., 2012).

Residual Risk with $\rho \equiv VaR_p$

- For simplicity, let

$$VaR_{p,\theta}(Y) = F^{-1}(p;\theta)$$

- Then positive residual risk implies an increased probability of default

$$VaR_{p,\theta}\left(Y - F^{-1}(p;\widehat{\theta})\right) \geq 0 \Leftrightarrow P_\theta\left(Y > F^{-1}(p;\widehat{\theta})\right) \geq 1-p$$

- VaR estimate is a *predictive limit* (Barndorff-Nielsen and Cox, 1996)

- Exponential distribution:

$$F(y;\theta) = 1 - \exp(-y/\theta),\ y > 0$$
$$P(Y > F^{-1}(p;\widehat{\theta})) = \left(1 - \frac{1}{n}\log(1-p)\right)^{-n}$$

Exponential example

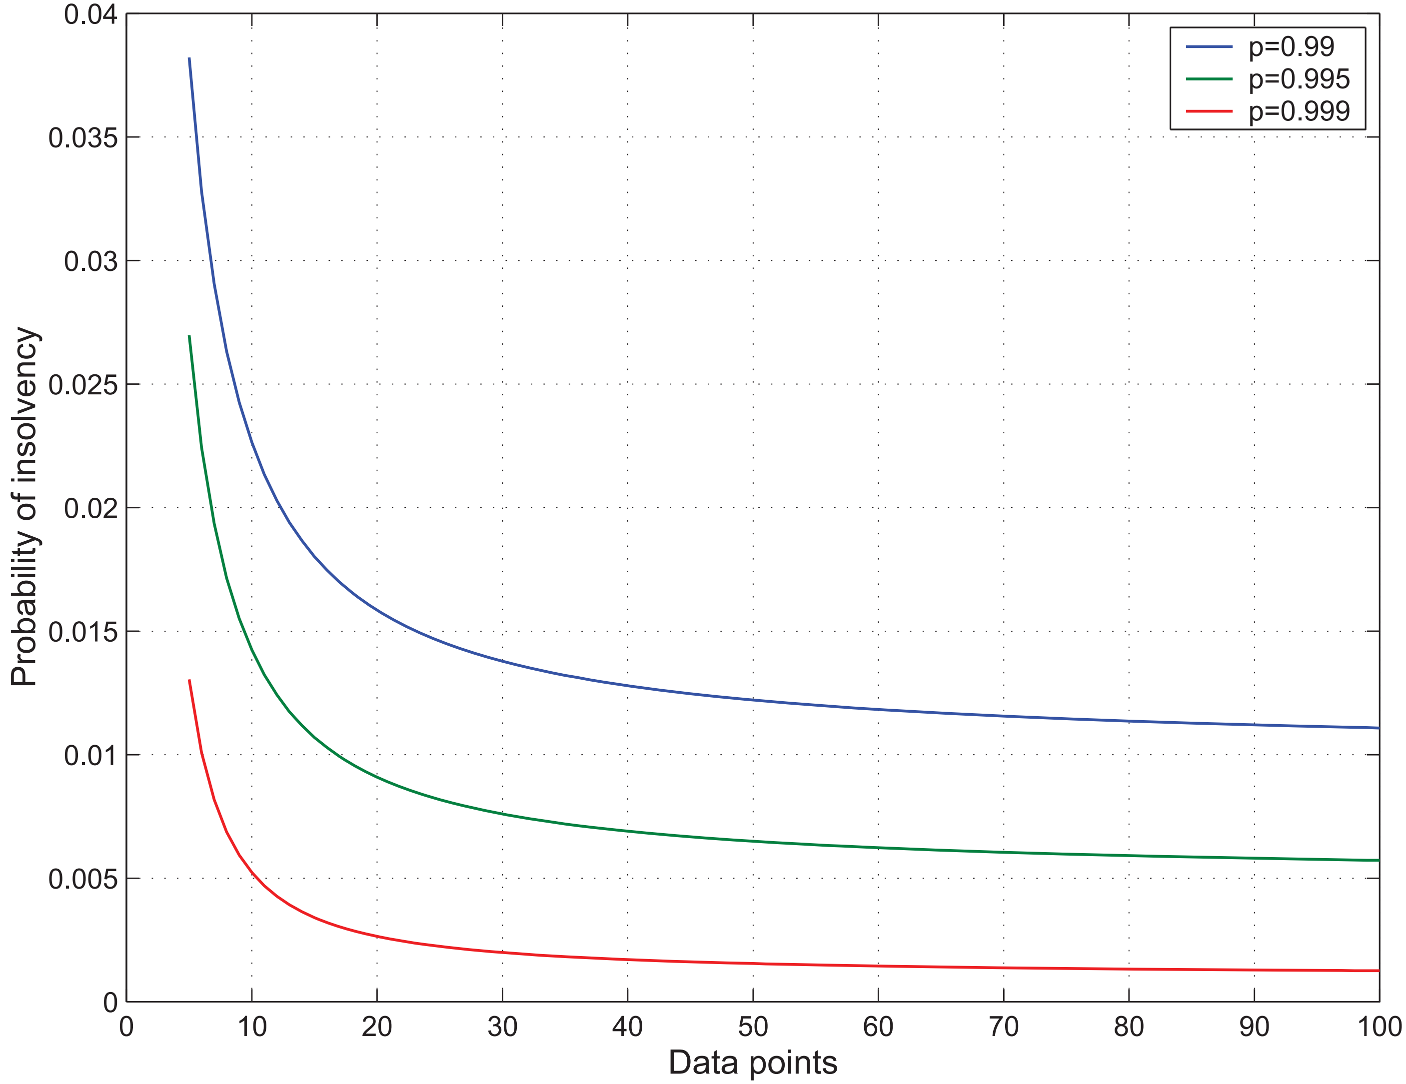

Residual Risk with $\rho \equiv VaR_p$

- Let Y be such that a known increasing transformation belongs to a location, scale, or location-scale family.
  - Normal, Log-Normal, Weibull, Log-Logistic, Exponential, Pareto...
- Then $P_\theta\left(Y > F^{-1}(p;\widehat{\theta})\right)$ does not depend on θ
  - Can find $p_n^* > p$ such that $P_\theta\left(Y > F^{-1}(p_n^*;\widehat{\theta})\right) = 1 - p$
- Let $\theta = (\mu, \sigma)$ and consider *probability matching prior* $\pi(\mu, \sigma) = 1/\sigma$ (Severini et al (2002); Gerrard & T. (2011))
  - Then $P_\theta(Y > \widehat{F}^{-1}(p|\mathbf{X})) = 1 - p$

Exponential example: $\widehat{\theta} = 100,\ n = 10$

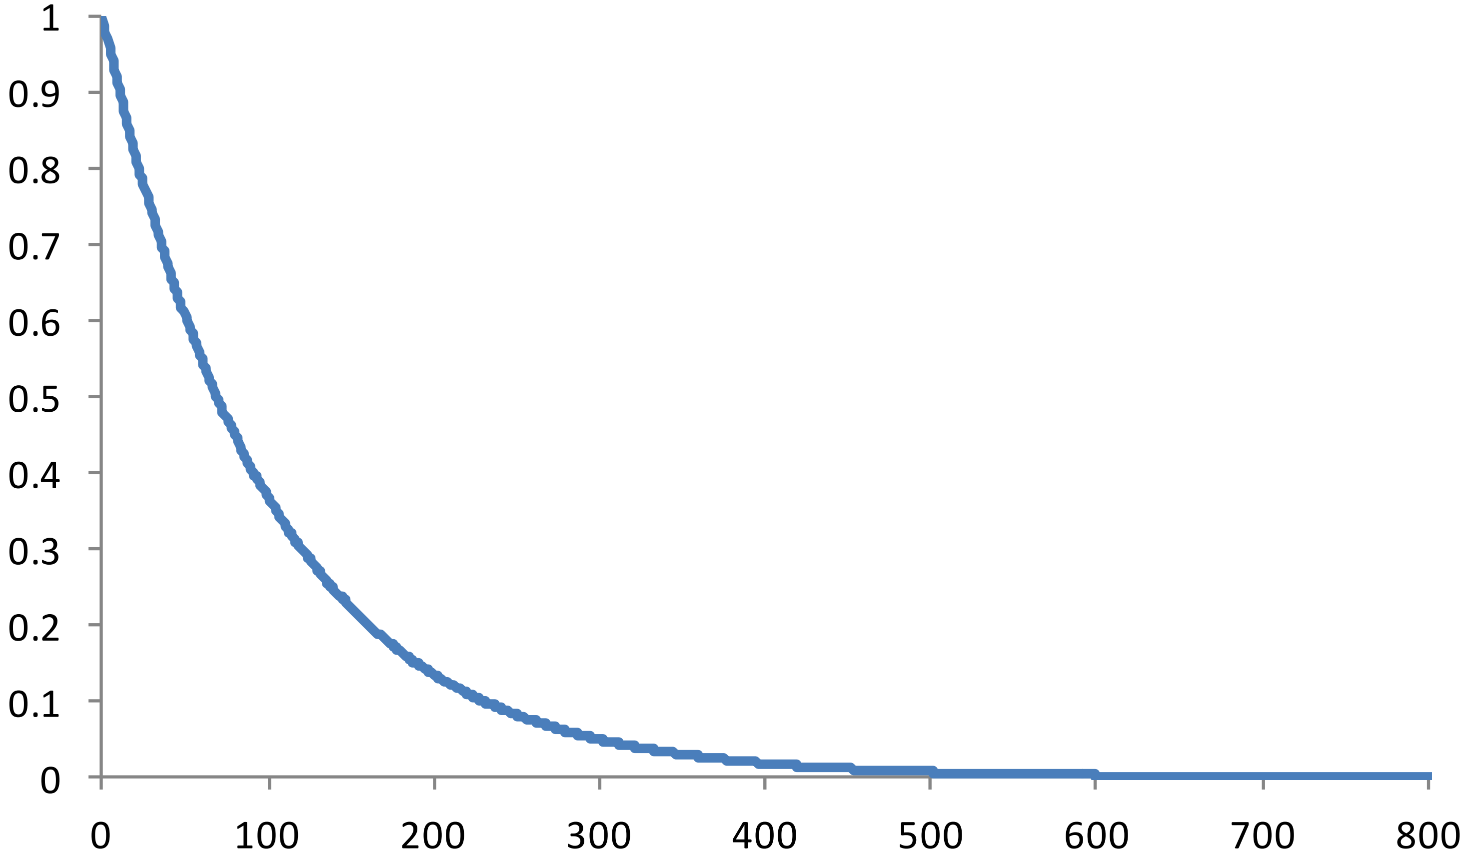

Exponential example: $\widehat{\theta} = 100,\ n = 10$

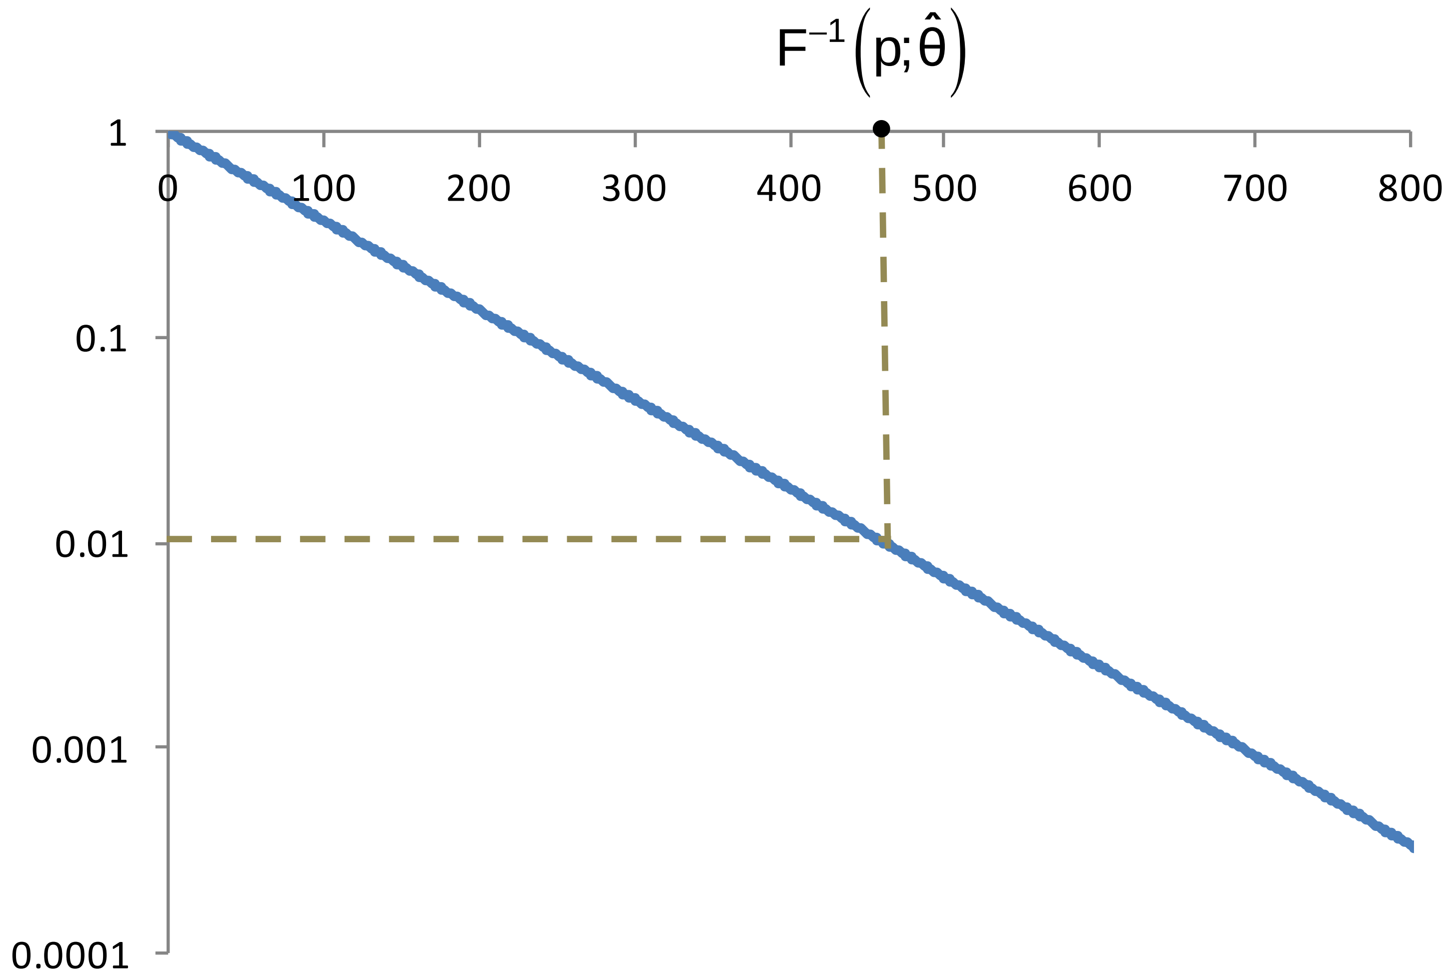

Exponential example: $\widehat{\theta} = 100,\ n = 10$

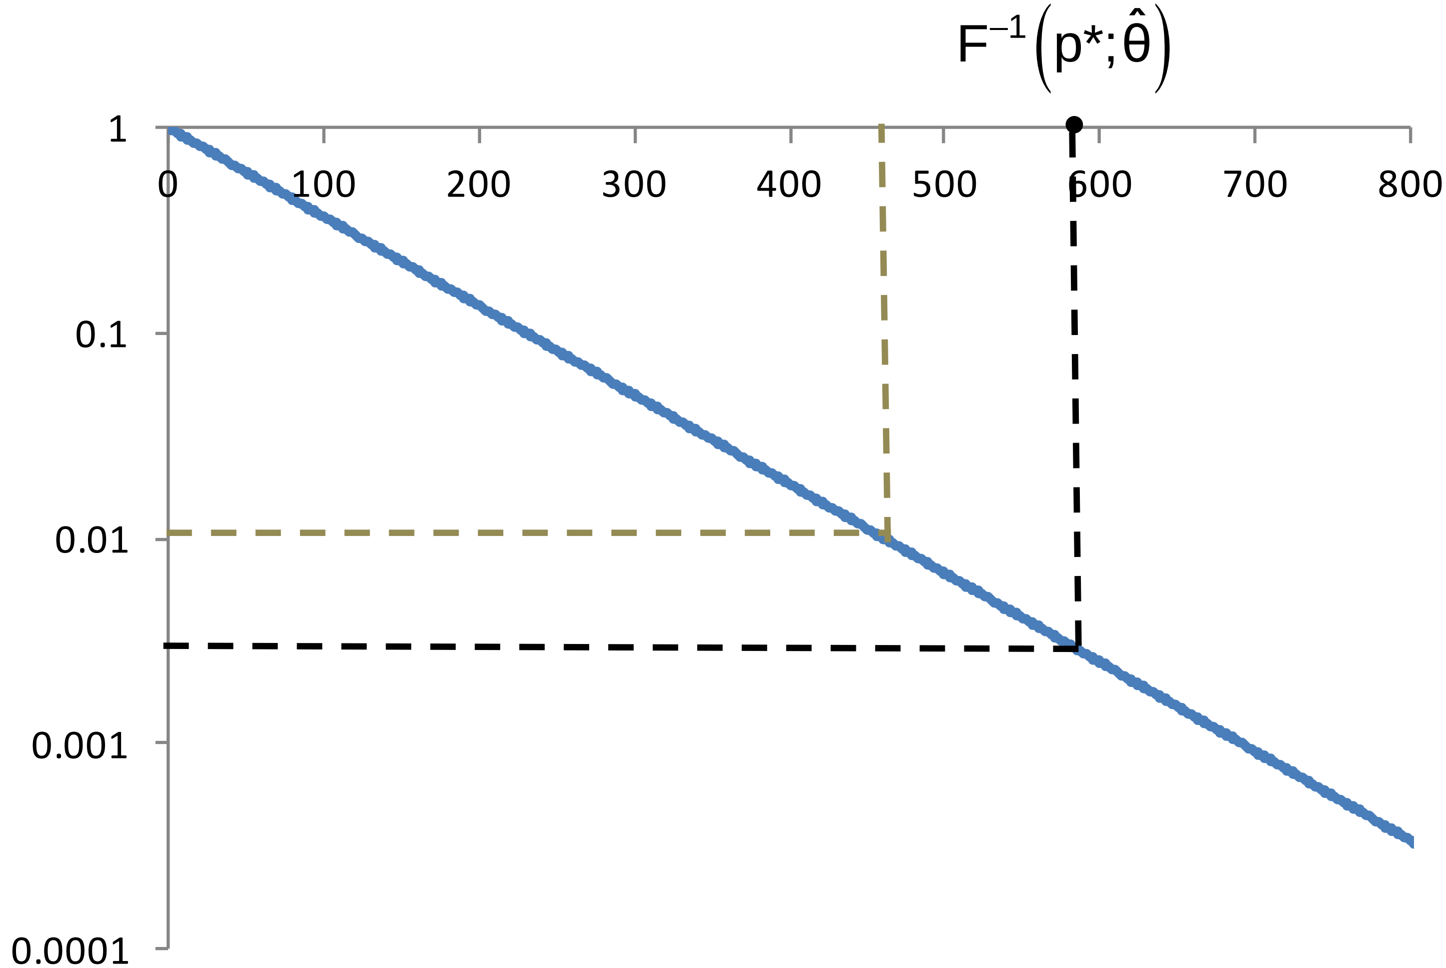

Exponential example: $\hat{\theta} = 100, \ n = 10$

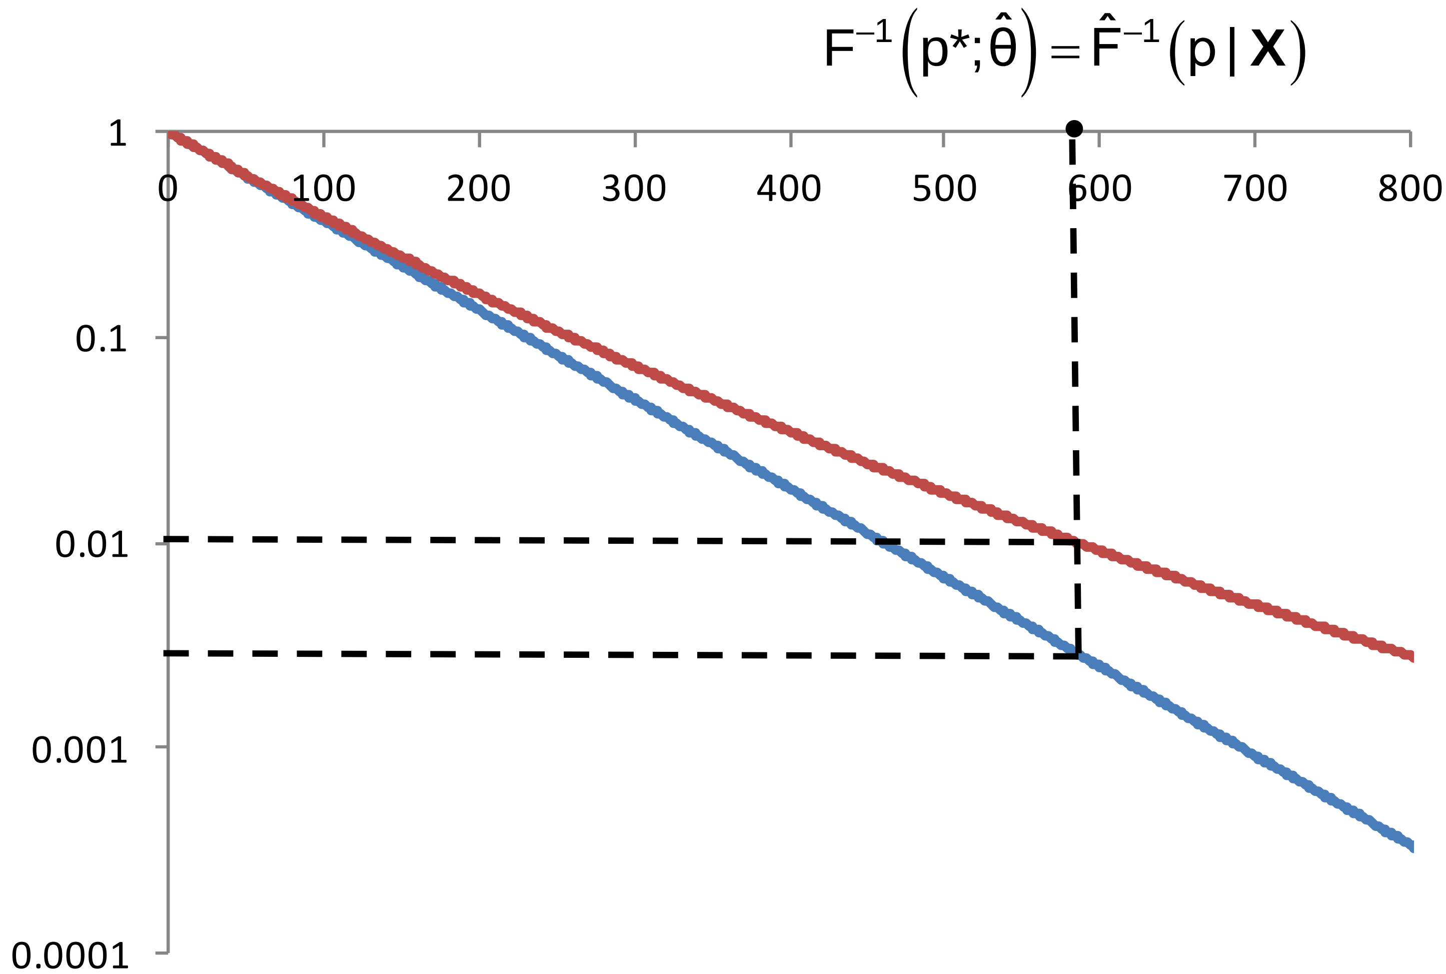

Beyond VaR

- Let ρ be some other risk measure eg $\rho \equiv TVaR_p$

- For location-scale families we find the following

  - One can construct another risk measure $\tilde{\rho}_n$ such that

    $$\rho_\theta\left(Y - \tilde{\rho}_{n,\widehat{\theta}}(Y)\right) = 0$$

  - Bayesian approach

    $$\rho_\theta\left(Y - \rho[\widehat{F}(\cdot|\mathbf{X})]\right) \approx 0$$

  - Parametric bootstrap (1st order)

    $$\rho_\theta\left(Y - \rho_{\widehat{\theta}}(Y) - \rho_{\widehat{\theta}}(Y - \rho_{\widehat{\theta}*}(Y))\right) \approx 0$$

    * Can improve by repeating the bootstrap without increasing computational time.

Normalised residual risk for exponential / $\rho \equiv TVaR$

MLE	**n=10**	**n=20**	**n=50**	**n=100**
p=0.95	0.2115	0.1179	0.0504	0.0261
p=0.99	0.2500	0.1442	0.0633	0.0333
p=0.995	0.2656	0.1552	0.0693	0.0355
Bayesian	**n=10**	**n=20**	**n=50**	**n=100**
p=0.95	-0.0181	-0.0088	-0.0024	-0.0015
p=0.99	-0.0126	-0.0051	-0.0034	-0.0010
p=0.995	-0.0108	-0.0053	-0.0031	0.0003
Bootstrap	**n=10**	**n=20**	**n=50**	**n=100**
p=0.95	0.0634	0.0202	0.0037	0.0010
p=0.99	0.0960	0.0332	0.0066	0.0017
p=0.995	0.1094	0.0389	0.0080	0.0021

Normalised residual risk for exponential / $\rho \equiv TVaR$

MLE	n=10	n=20	n=50	n=100
p=0.95	0.2115	0.1179	0.0504	0.0261
p=0.99	0.2500	0.1442	0.0633	0.0333
p=0.995	0.2656	0.1552	0.0693	0.0355
Bayesian	**n=10**	**n=20**	**n=50**	**n=100**
p=0.95	-0.0181	-0.0088	-0.0024	-0.0015
p=0.99	-0.0126	-0.0051	-0.0034	-0.0010
p=0.995	-0.0108	-0.0053	-0.0031	0.0003
Bootstrap2	**n=10**	**n=20**	**n=50**	**n=100**
p=0.95	0.0298	0.0053	0.0004	0.0001
p=0.99	0.0679	0.0137	0.0012	0.0002
p=0.995	0.0900	0.0193	0.0017	0.0002

Something about time consistency

- Let θ_0 be initial guess of parameter

  $$\rho_{\theta_0}\Big(Y - \tilde{\rho}_{n,\widehat{\theta}}(Y)\Big) = 0$$

- Implication:

  $$\tilde{\rho}_{n,\widehat{\theta}}(Y) \leq 0 \; a.s. \implies \rho_{\theta_0}(Y) \leq 0$$
  $$\tilde{\rho}_{n,\widehat{\theta}}(Y) \geq 0 \; a.s. \implies \rho_{\theta_0}(Y) \geq 0$$

- *Sequentially consistent* risk measurement (Roorda and Schumacher (2007); Bignozzi and T. (2012))

Something about heavy tails

- If ρ is *coherent*, then finite means are required
- Let $Y \sim LN(\mu, \sigma^2)$
  - Predictive distribution of Y is Log-t $\rightarrow$ infinite mean!
- Let $Y \sim Pareto(\theta)$ such that $P_\theta(Y > y) = y^{-\theta},\ y \geq 1,\ \theta > 1$.
  - $P_\theta(\widehat{\theta} < 1) > 0 \implies P_\theta(\rho_{\widehat{\theta}}(Y) < \infty) < 1$
- Blame the statistics or the risk measure?
- Could follow Cont et al (2010) and use

$$\rho_\theta(Y) = \frac{1}{p_2 - p_1} \int_{p_1}^{p_2} VaR_{\alpha,\theta}(Y) d\alpha$$

Something about pricing

- Theoretical price of layer l in excess of d

$$\pi[F(\cdot;\theta)] = \int_d^{d+l} (1 - F(y;\theta))d\theta$$

- What should we use as premium?
  - MLE $\pi[F(\cdot;\widehat{\theta})]$ v predictive distribution $\pi[\widehat{F}(\cdot|\mathbf{X})]$

- Let portfolio consist of many independent policies, with premiums calculated from different data
  - Consider $\sum_j \left(Y_j - \pi[\widehat{F}(\cdot|\mathbf{X})]\right)$
  - With increasing portfolio size, premium bias becomes dominant
  - Estimation volatility *diversifies away*

Relative bias of MLE of exponential tail function

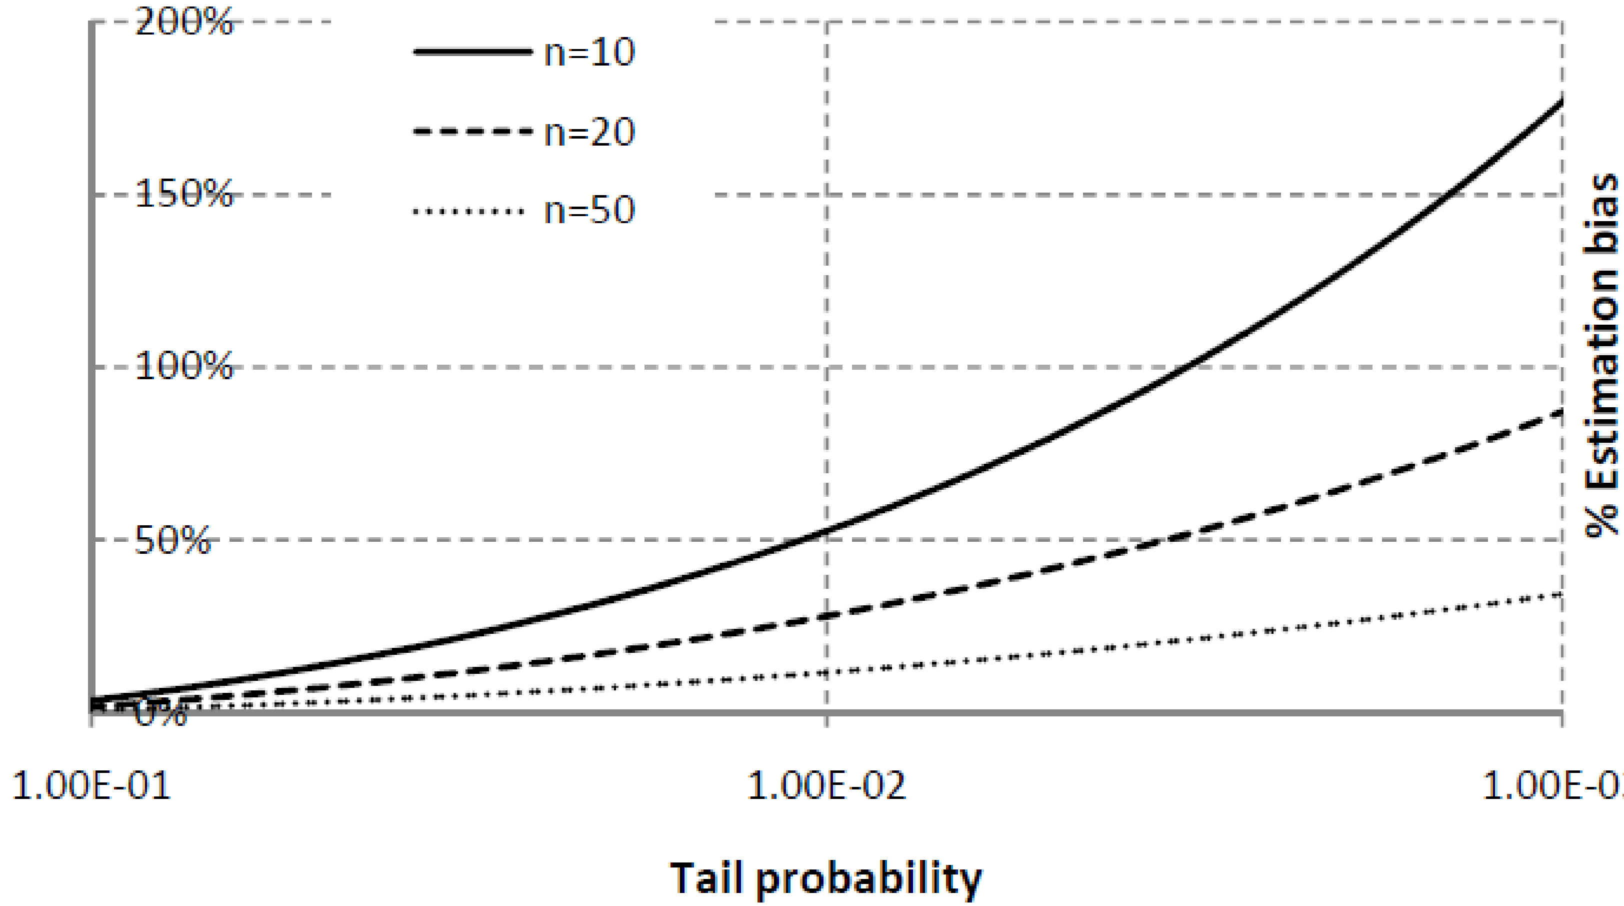

Exact and approximate bias (Landsman and T., 2012)

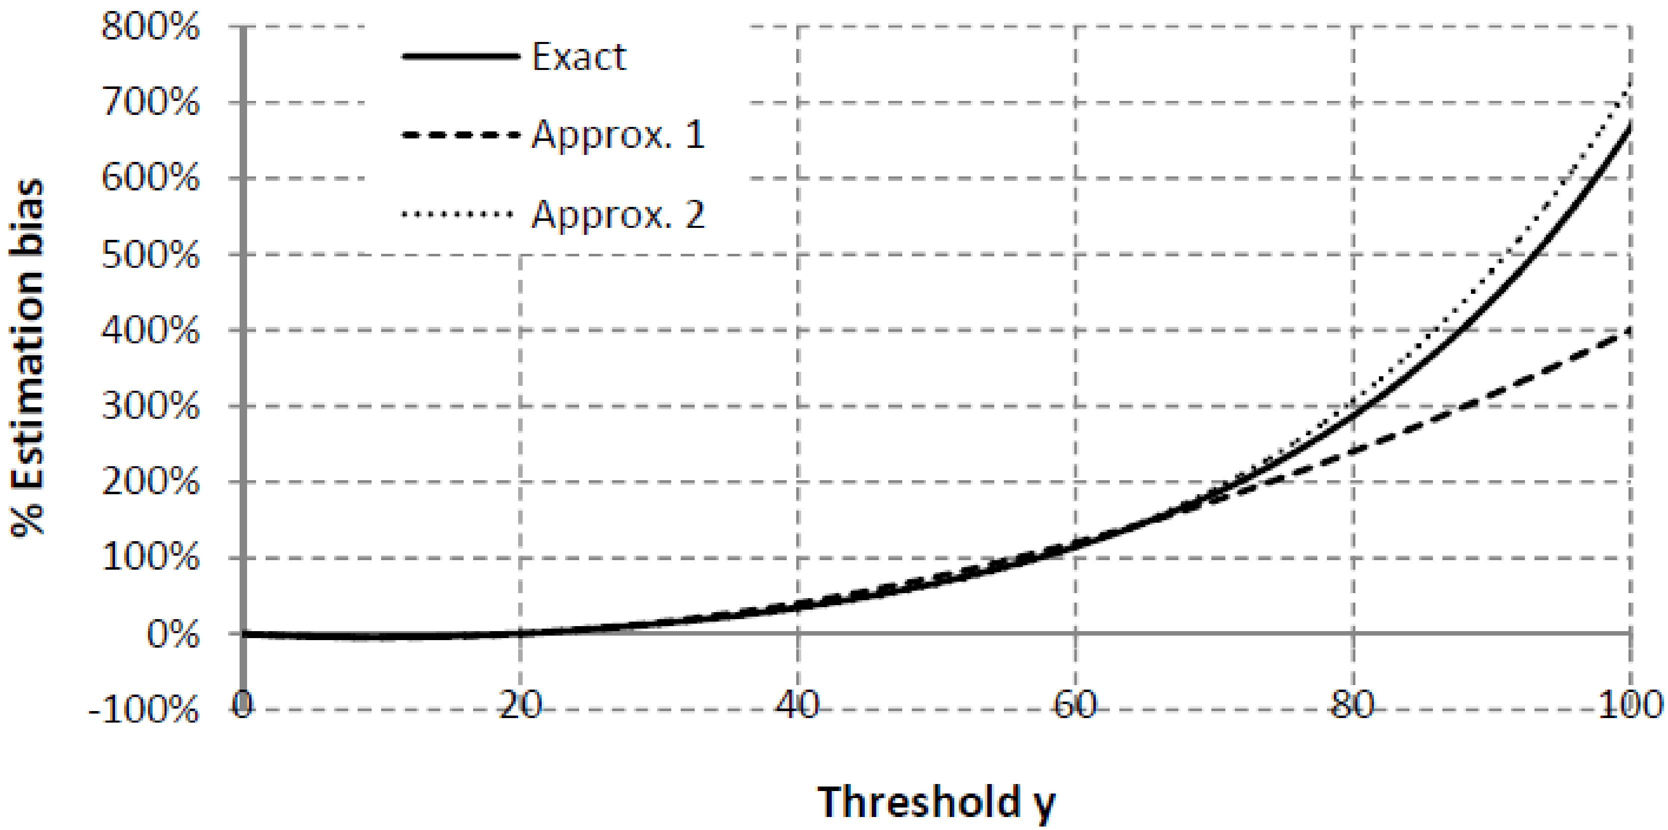

MLE v predictive bootstrap distribution v bias-correction

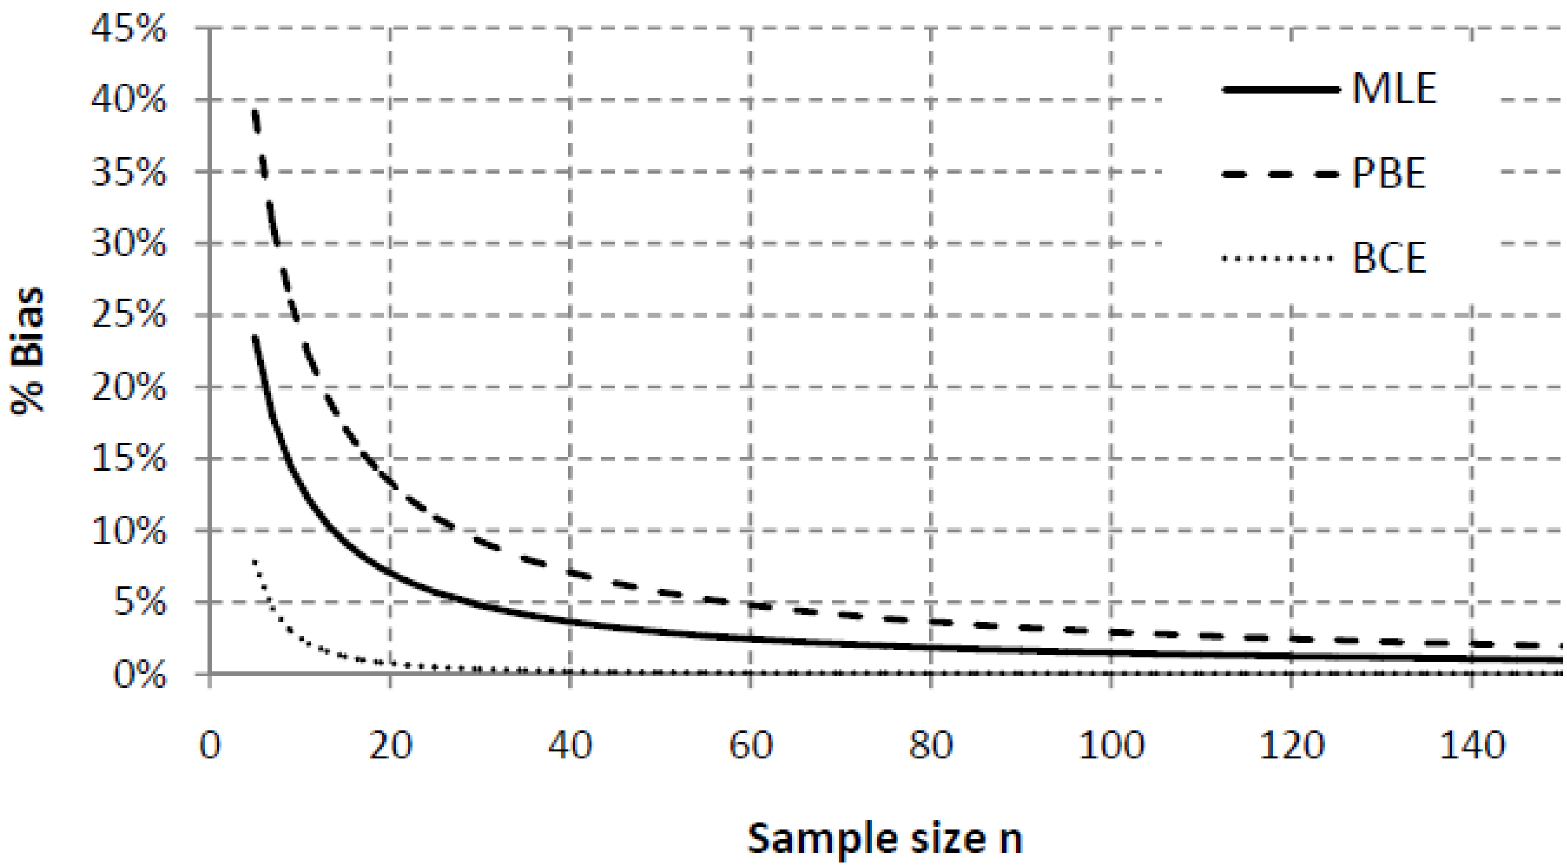

References

- Cairns (2000) A discussion of parameter and model uncertainty in insurance. *Insurance: Mathematics and Economics*.

- Barndorff-Nielsen and Cox (1996) Prediction and asymptotics. *Bernoulli*.

- Severini, Mukerjee and Ghosh (2002) On exact probability matching property of right-invariant priors. *Biometrica*.

- Gerrard and Tsanakas (2011) Failure Probability Under Parameter Uncertainty. *Risk Analysis: An International Journal*.

- Bignozzi and Tsanakas (2012) Parameter uncertainty and residual risk.

- R. Cont and R. Deguest and G. Scandolo (2010) Robustness and sensitivity analysis of risk measurement procedures. *Quantitative Finance*.

- Roorda and Schumacher (2007) Time consistency conditions for acceptability measures with an application to tail value at risk. *Insurance: Mathematics and Economics*.

- Bignozzi and Tsanakas (2012) Characterisation and construction of sequentially consistent risk measures.

- Landsman and Tsanakas (2012) Parameter uncertainty in exponential family tail estimation. *ASTIN Bulletin*