

MERCER OLIVER WYMAN

October 22-23, 2004

Comments on De Fontnouvelle, Rosengren, and Jordan
"Implications of Alternative Operational Risk Modelling
Techniques"

NBER Risks of Financial Institutions Conference
Woodstock, Vermont

Confidentiality

Our clients' industries are extremely competitive. The confidentiality of companies' plans and data is obviously critical. Mercer Oliver Wyman will protect the confidentiality of all such client information.

Similarly, management consulting is a competitive business. We view our approaches and insights as proprietary and therefore look to our clients to protect Mercer Oliver Wyman interests in our presentations, methodologies and analytical techniques. Under no circumstances should this material be shared with any third party without the written consent of Mercer Oliver Wyman.

Copyright © 2004 Mercer Oliver Wyman

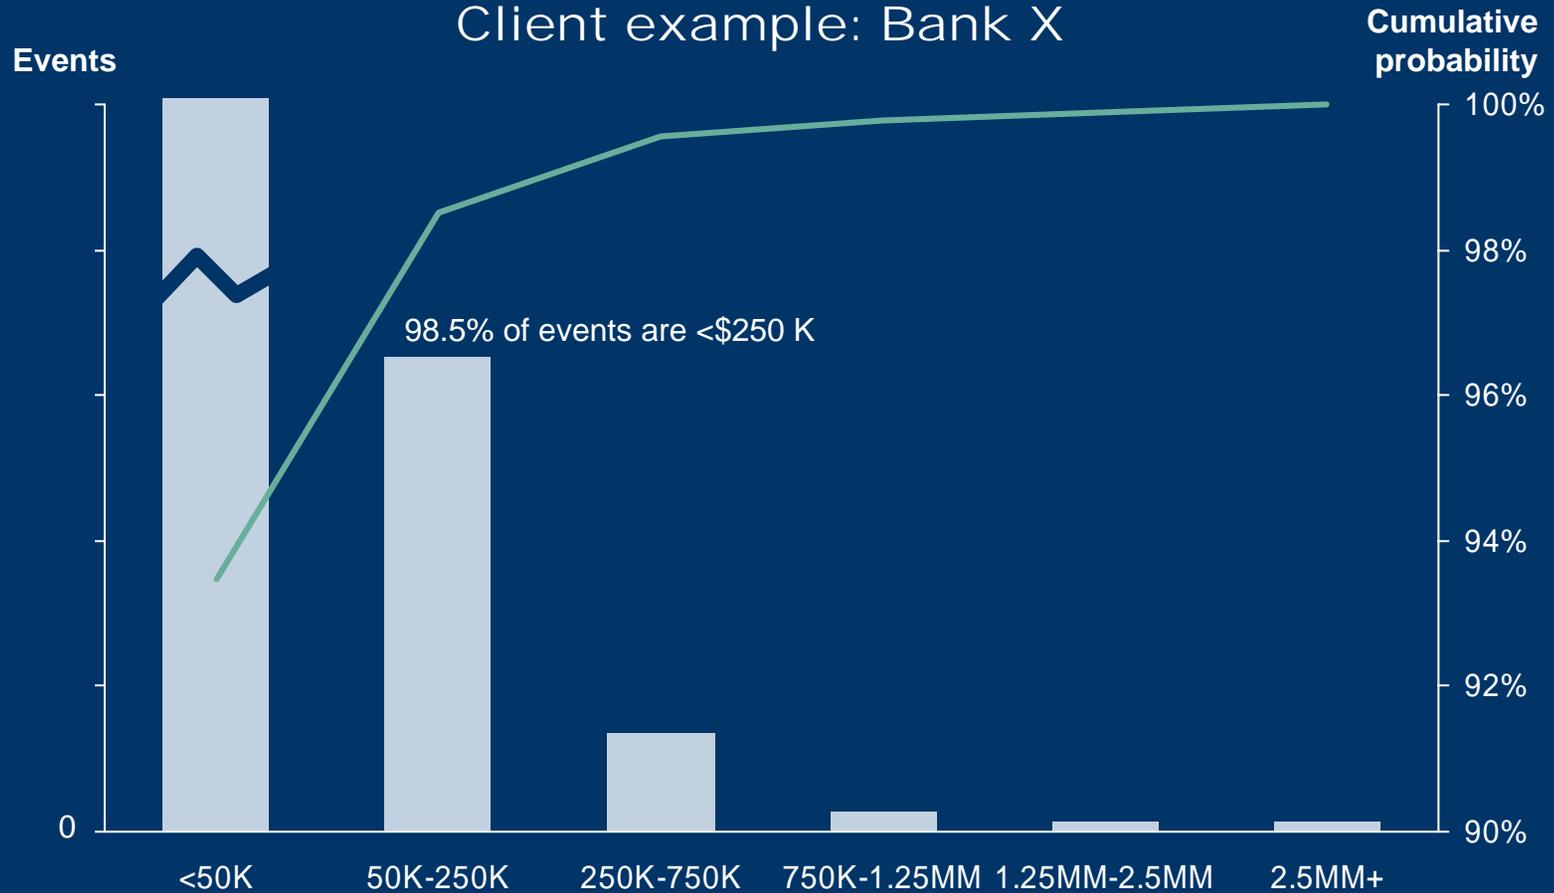
Key questions

1. Are operational losses best characterized by a thin-tailed (e.g., lognormal) or fat-tailed (e.g., Pareto) distribution?
2. Do the estimates of operational losses across individual banks converge?
3. Do the results provide a reasonable basis for allocating capital?

The Problem (1): Paucity of internal data on large losses

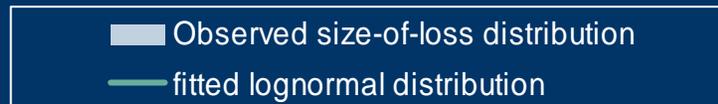
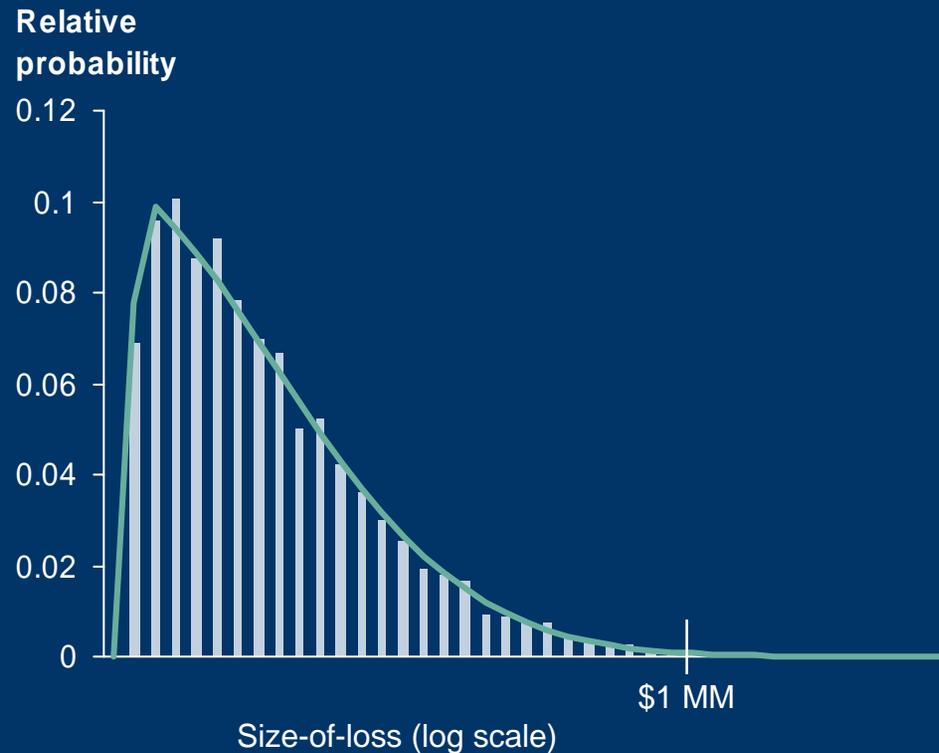
5-years cumulative operational losses

Client example: Bank X



The Problem (2): Internal data appears to fit a lognormal distribution

Client example: Bank X
Empirical vs. log normal



Modelled quantiles scaled
to bank assets at 99.9%

From De Fontnouvelle, Table 5

Poisson A – Lognormal	0.056%
Poisson A – Empirical	0.053%
Poisson B – Lognormal	0.101%
Poisson B – Empirical	0.093%
Neg Binomial – Lognormal	0.198%
Neg Binomial – Empirical	0.202%

Counter-evidence

1. For capital purposes, not concerned with mode of distribution, but with tail at $\geq 99.9\%$
2. Experience in modeling extreme events (e.g. CAT losses) suggests tails can behave very differently from the body of distribution
3. Profiling of operational loss data across six banks (tail plots, average excess plots, chi-square) indicates that fat-tailed distributions generally outperform than thin-tailed
4. Using HKKP estimation, Pareto distribution appears to provide a strong fit for the six banks
5. Results are consistent with earlier work by de Fontnouvelle et al examining publicly-reported operational losses

Proposed solution

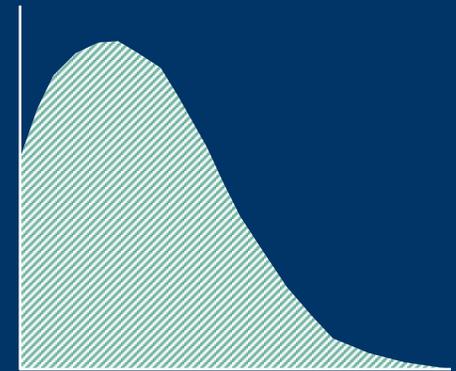
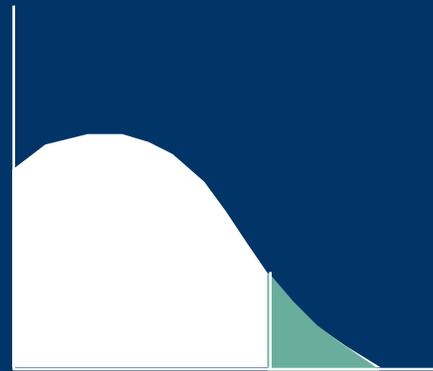
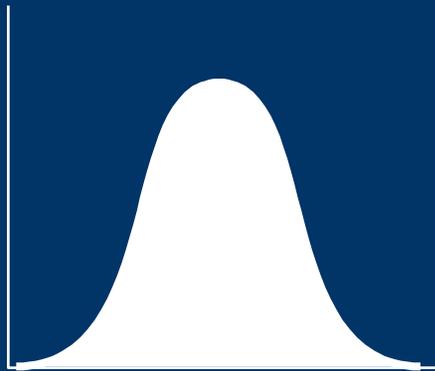
Frequency

x

Severity

=

Loss

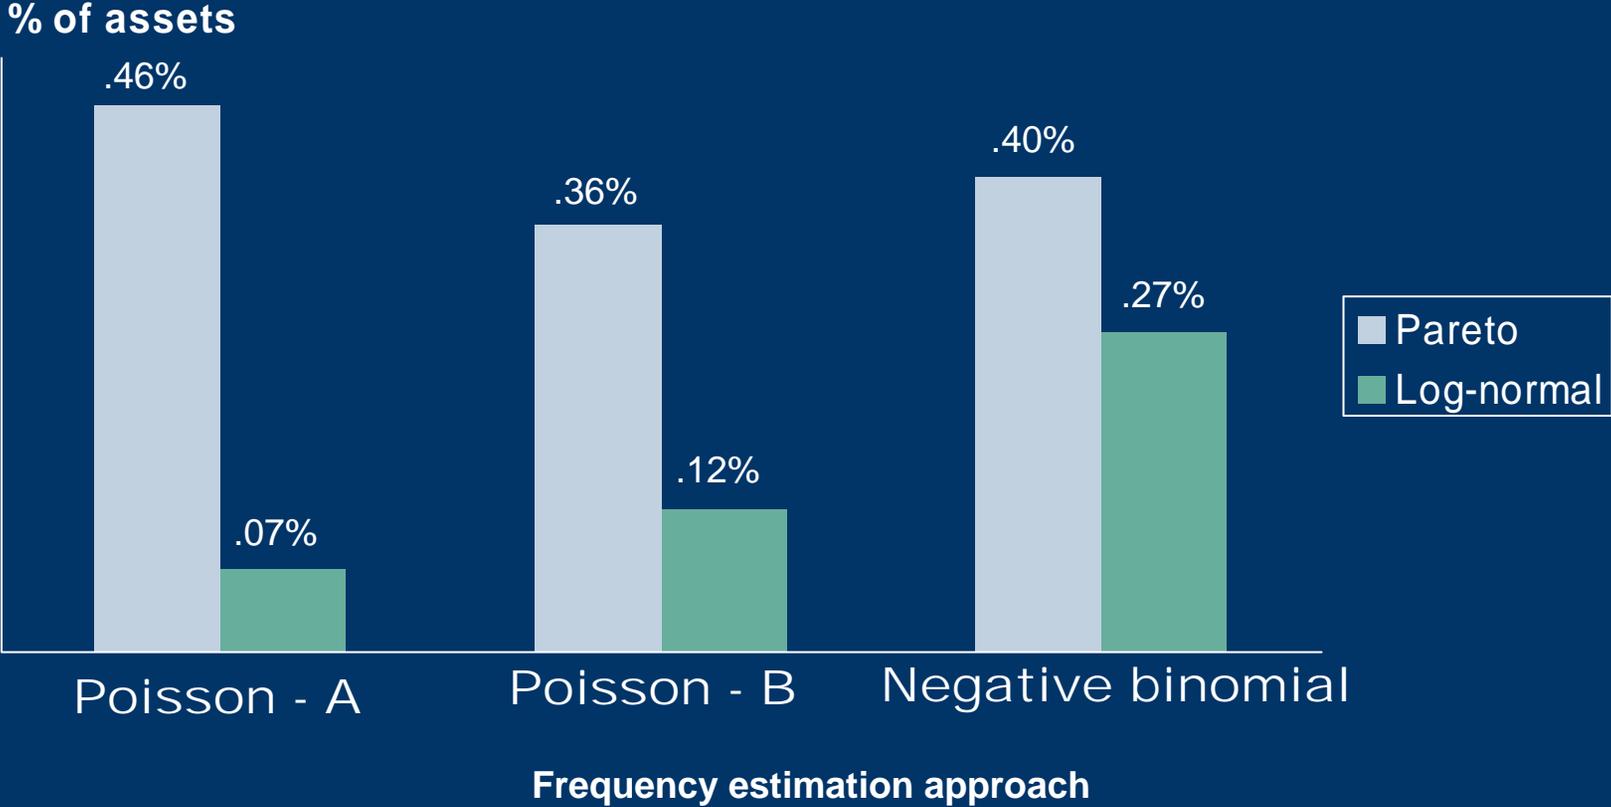


■ Based on internal loss data

■ Based on external data modeling

Implications of thin-tailed vs fat-tailed distributions

Modeled capital at 99.9% level



Notes/Source: 8 point, in same place on every page if possible

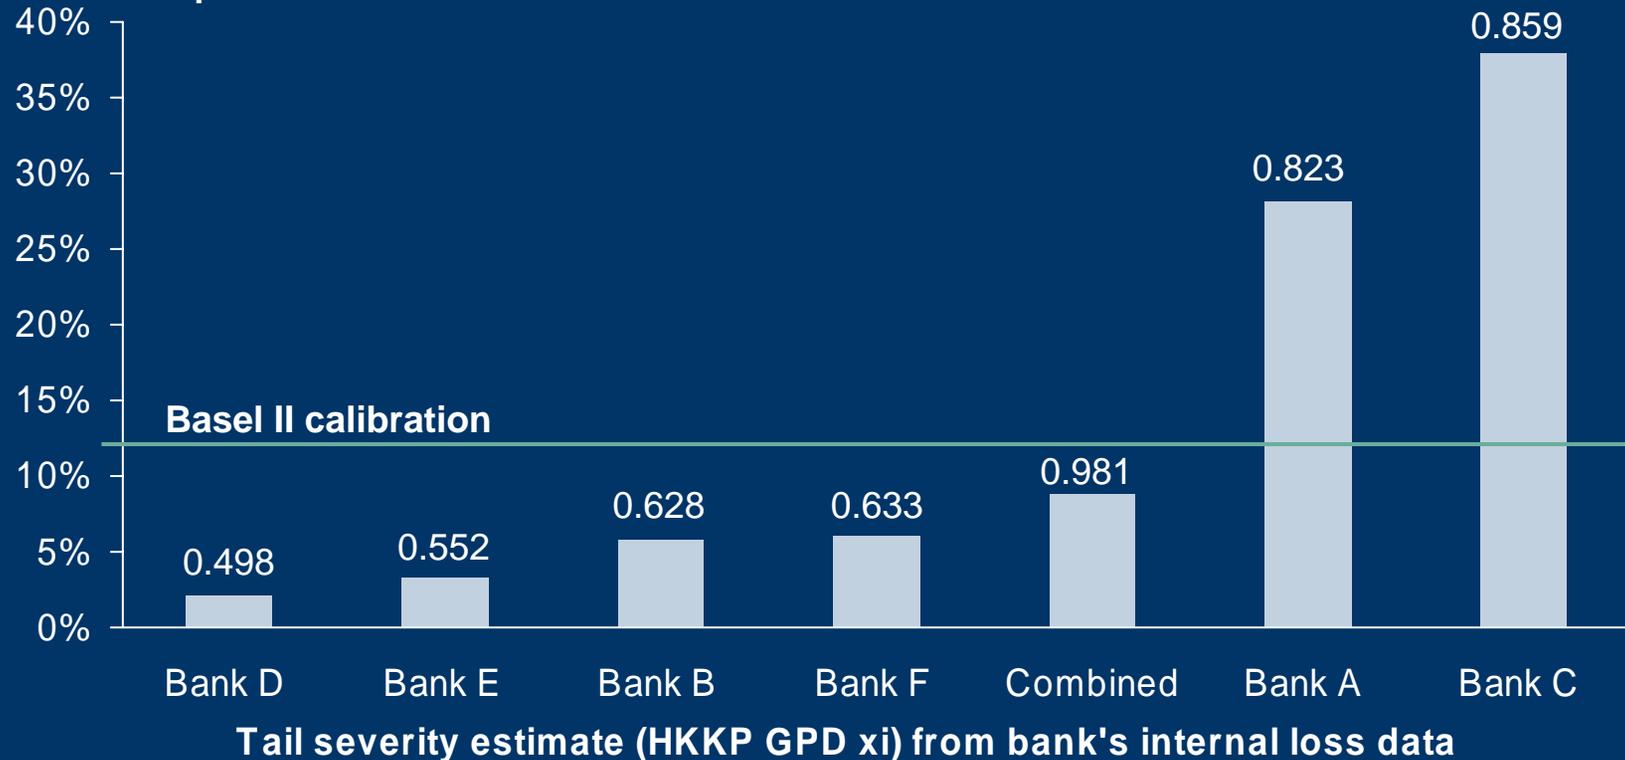
Do estimates across banks converge?

Bank	Exceedances	K
A	180	0.82
B	80	0.62
C	30	0.85
D	50	0.49
E	200	0.55
F	50	0.63
All	140	0.68
External	\$1 MM	0.65

Do the results provide a reasonable basis for allocating capital?

Impact of different shape parameters on Bank X

Ratio of OpRisk to total capital



Conclusions

1. Are operational losses thin-tailed or fat-tailed?
 - Convincing case for fat-tailed distributions (GDP), and for modeling op losses using “peaks over threshold” approach
2. Do estimates converge across individual banks?
 - Some evidence of convergence between internal and external data sets
 - But seemingly wide disparity in estimates at individual bank level (based on 1-year data)
3. Do results provide reasonable basis for setting capital?
 - Depends on shape parameter (10:1 range of results)
 - Impossible to tell based on internal analysis alone