Reverse Survivorship Bias

Citation format:
Linnainmaa, Juhani T., 2013, Reverse Survivorship Bias,
Journal of Finance 68, 789–813

Fund managers' stock-picking abilities

Do any mutual fund managers have enough skill to cover the costs they impose on their investors?

- Equilibrium accounting: the average actively managed dollar *must* have a negative alpha
- Heterogeneity?
- Barras, Scaillet, and Wermers (2010):

  "While these papers are useful in uncovering whether, on the margin, outperforming mutual funds exist, they are not particularly informative regarding their prevalence in the entire fund population. For instance, it is natural to wonder how many fund managers possess true stockpicking skills [. . .]"

(Direct) Survivorship bias

- Early fund databases lacked data on dead funds
- If funds disappear following poor performance, conditioning on survival leads to an upward bias in estimated alphas
  - Carhart, Carpenter, Lynch, and Musto (2002): bias as large as 1% in samples longer than 15 years
- Starting from Elton, Gruber, and Blake (1996) and Carhart (1997) databases are survivorship-bias free
  - Include both surviving and dead funds

Example 1: Learning about alphas

- Suppose investors do not know funds' true alphas
  - Learn over time from risk-adjusted returns
- Suppose a fund disappears when the posterior mean about its alpha falls below some fixed threshold
- Bayes' rule: Posterior mean between the prior and signal
- Key observation: If the threshold is -1%, it can only be crossed if the in-sample alpha estimate ("signal") is *strictly less* than -1%

Reverse Survivorship Bias

The gap between the expected alpha and in-sample alpha estimate is the **reverse survivorship bias**

Example 1: Learning about alphas

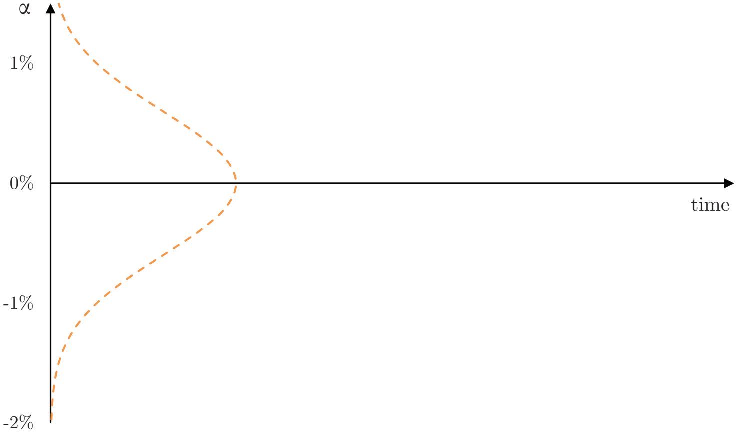

Example 1: Learning about alphas

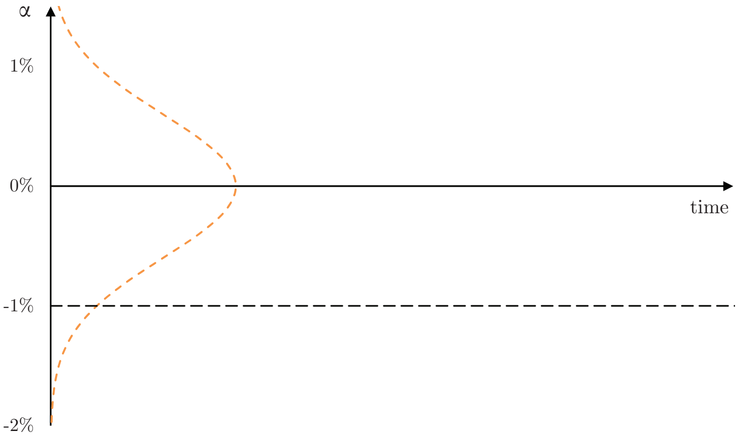

Example 1: Learning about alphas

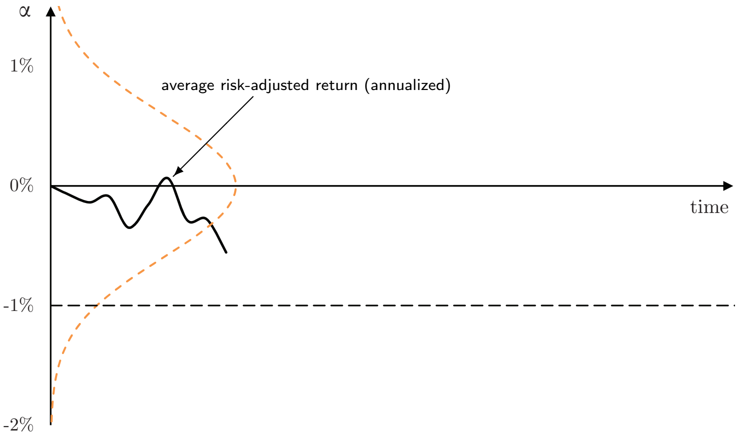

Example 1: Learning about alphas

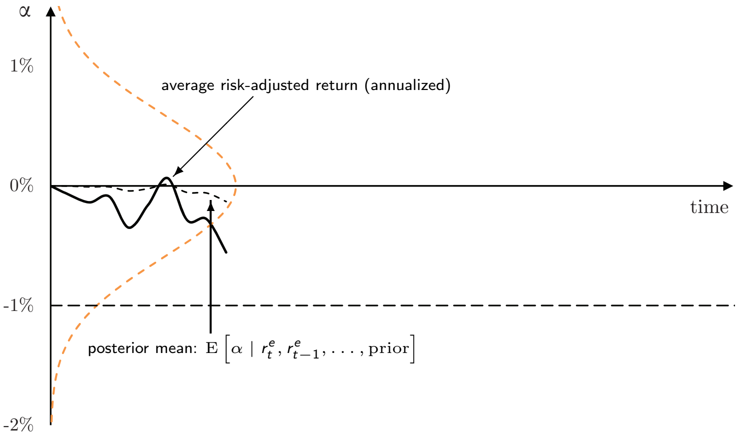

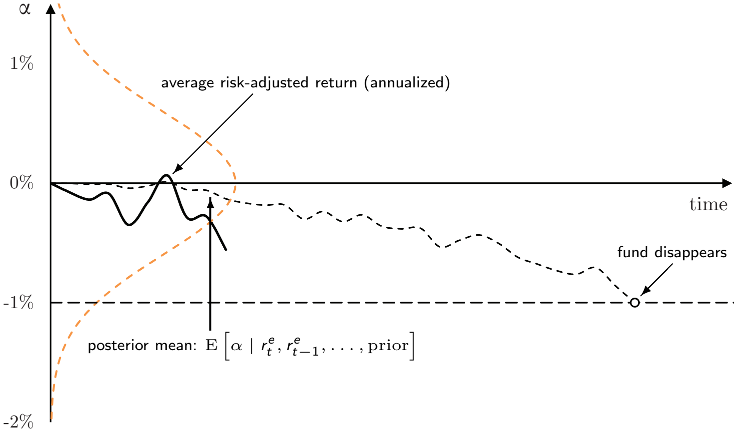
α
1%
0%
-1%
-2%
time
average risk-adjusted return (annualized)
fund disappears
posterior mean: $\mathrm{E}\left[\alpha \mid r_t^e, r_{t-1}^e, \ldots, \text{prior}\right]$

Example 1: Learning about alphas

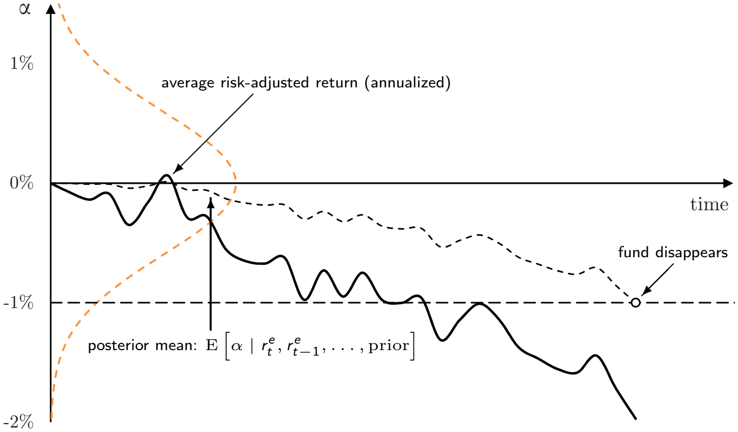

Example 1: Learning about alphas

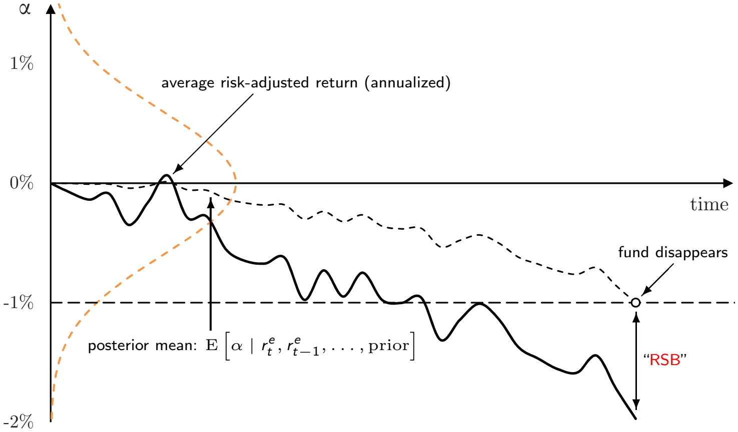

Example 2: No learning

- This bias does not require learning
- The simplest example:
  1. All mutual funds identical
  2. Two periods: return 10% or -10%, equal probabilities
  3. A fund disappears if it posts a low return
- What kind of data do these rules generate?

Fraction of Funds	Sequence	Mean Return
50%	Down	-10%
25%	Up-Down	0%
25%	Up-Up	10%

→ The expected realized alpha is -2.5%

If mutual funds "stop" after low risk-adjusted returns, their estimated alphas are biased downwards relative to true alphas

Example 2: No learning

- This bias does not require learning
- The simplest example:
  1. All mutual funds identical
  2. Two periods: return 10% or −10%, equal probabilities
  3. A fund disappears if it posts a low return
- What kind of data do these rules generate?

Fraction of Funds	Sequence	Mean Return
50%	Down	−10%
25%	Up-Down	0%
25%	Up-Up	10%

→ The expected realized alpha is −2.5%

If mutual funds "stop" after low risk-adjusted returns, their estimated alphas are biased downwards relative to true alphas

Formal argument

Assumptions

1. Fund generates risk-adjusted returns, $r_t^e = \alpha + \tilde{\varepsilon}_t$, where α is constant and $\tilde{\varepsilon}_t$ mean zero and i.i.d.
2. Fund survives for a random $\tilde{T}$ number of months

- The optional stopping-time theorem: $\mathrm{E}\left[\sum_{t=1}^{\tilde{T}} \tilde{\varepsilon}_t\right] = 0$
- The expected average risk-adjusted return is then

$$\mathrm{E}\left[\frac{1}{\tilde{T}}\sum_{t=1}^{\tilde{T}} \tilde{r}_t^e\right] = \alpha + \mathrm{cov}\left(\frac{1}{\tilde{T}}, \sum_{t=1}^{\tilde{T}} \tilde{\varepsilon}_t\right)$$

$\Rightarrow$ **If survival correlates positively with risk-adjusted returns,** $\mathbf{E}\left[\frac{1}{\tilde{T}}\sum_{t=1}^{\tilde{T}} \tilde{r}_t^e\right] < \alpha$

Reverse survivorship bias (I)

- Funds often disappear following poor performance. . .
  - But performance can be poor not only because the true alpha is low but also because of negative idiosyncratic shocks
  - The market cannot perfectly disentangle luck from skill

$\Rightarrow$ estimated alphas too low for dead funds

- Important bias for questions about fund managers' abilities:
  1. How many managers have enough skill to cover the costs?
  2. How easy or difficult to identify skilled managers from past returns alone?

  etc.

Reverse survivorship bias (II)

- This bias does not matter when:
  - We study returns available to mutual fund investors...
  - ...or if we only care about the mean
- A long tradition of estimating alphas fund by fund, going back to Jensen (1968)
- Many recent papers take a fund-level focus:
  - Kosowski, Timmermann, Wermers, and White (2006): sizable minority of managers pick stocks well enough to more than cover their costs
  - Barras, Scaillet, and Wermers (2010): (reliably) negative-alpha funds (24%) outnumber positive-alpha funds (0.6%)
  - Fama and French (2010): few funds produce benchmark adjusted expected returns sufficient to cover their costs

Reverse survivorship bias (II)

- This bias does not matter when:
  - We study returns available to mutual fund investors...
  - ...or if we only care about the mean
- A long tradition of estimating alphas fund by fund, going back to Jensen (1968)
- Many recent papers take a fund-level focus:
  - Kosowski, Timmermann, Wermers, and White (2006): sizable minority of managers pick stocks well enough to more than cover their costs
  - Barras, Scaillet, and Wermers (2010): (reliably) negative-alpha funds (24%) outnumber positive-alpha funds (0.6%)
  - Fama and French (2010): few funds produce benchmark adjusted expected returns sufficient to cover their costs

Road Map

1. CRSP mutual fund data
2. Alpha estimates for surviving and dead funds
3. Bias in the mean and simulation evidence
4. Structural estimation
5. Remedies

CRSP Survivorship-Bias Free Mutual Fund Database

- Fama and French (2010) sample construction:
  1. U.S. equity mutual funds
  2. Use data on funds that start in 1984 or later
  3. Exclude funds with fewer than 9 months of data
  4. Combine multiple share classes of the same fund
  5. Leave out returns for as long as AUM $<$ $5M (2006 dollars)

     $\rightarrow$ Incubation bias screen (Evans 2010)
- Return data from January 1984 through September 2010, 1,853 funds

Alpha estimates conditional on survival

Model	Survive	Fund Age in Years: 1	3	5	7	10
CAPM	Yes	−0.62	−1.16	0.15	0.54	0.75
		(−2.11)	(1.72)	(0.98)	(4.58)	(7.66)
	No	−5.92	−5.42	−4.37	−2.31	−2.14
		(−3.77)	(−6.11)	(−6.25)	(−4.96)	(−6.04)
Four-Factor Model	Yes	−0.09	−0.15	0.19	0.10	0.03
		(−0.37)	(−0.91)	(1.45)	(0.99)	(0.34)
	No	−4.57	−3.49	−4.03	−1.97	−2.04
		(−4.41)	(−5.44)	(−6.78)	(−5.01)	(−5.62)
Survival Rate		0.995	0.943	0.868	0.789	0.681

Cumulative risk-adjusted returns conditional on survival

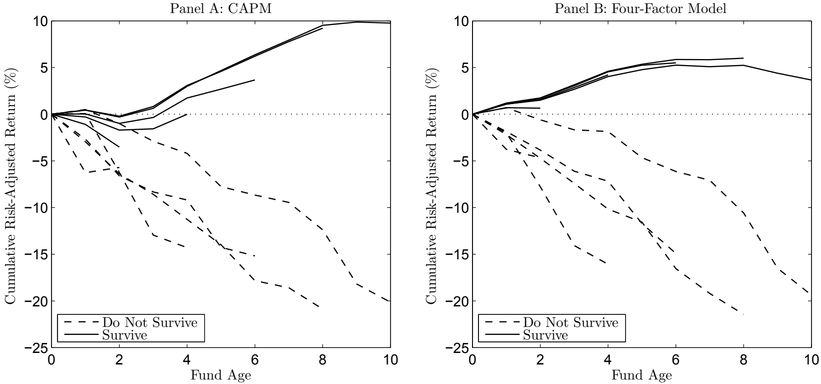

Cox proportional hazards model for fund survival

	Covariate: $\hat{\alpha}$			Covariate: $t(\hat{\alpha})$		
	CAPM	FF	FF+M	CAPM	FF	FF+M
Hazard ratio	0.913	0.900	0.891	0.642	0.713	0.707
SE	[0.005]	[0.006]	[0.006]	[0.019]	[0.016]	[0.018]
Pseudo R^2	0.019	0.020	0.022	0.024	0.015	0.015

Bias in the Mean

and

Simulation Evidence

Bias in the mean of the alpha distribution (I)

- Assume fund returns conform to a factor model:

$$r_{j,t} - r_{f,t} = \alpha_j + F_t \beta_j + \varepsilon_{j,t}$$

- Two methods for estimating the average alpha:
  1. Run fund-by-fund regressions, compute $\frac{1}{N}\sum_{j=1}^{N} \hat{\alpha}_j$
  2. Compute average fund return for each month and estimate *one* time-series regression:

$$\frac{1}{N}\sum_{j=1}^{N} r_{j,t} - r_{f,t} = \frac{1}{N}\sum_{j=1}^{N} \alpha_j + F_t \frac{1}{N}\sum_{j=1}^{N} \beta_j + \frac{1}{N}\sum_{j=1}^{N} \varepsilon_{j,t}$$

Bias in the mean of the alpha distribution (II)

1. Average alphas from fund-by-fund regressions:
   - CAPM: -0.69% ($t = -6.4$)
   - Three-factor model: -1.10% ($t = -12.0$)
   - Four-factor model: -1.14% ($t = -13.2$)
2. Average-return time-series regressions:
   - CAPM: -0.16% ($t = -0.3$)
   - Three-factor model: -0.51% ($t = -1.4$)
   - Four-factor model: -0.53% ($t = -1.4$)

Simulation evidence: Assumptions (I)

1. Alphas drawn from a true distribution $N(0, 1.25\%)$
2. Investors extract risk-adjusted returns from raw returns:

$$\tilde{r}^e_{j,t} = \alpha_j + \tilde{\varepsilon}_{j,t}$$

3. $\tilde{\varepsilon}_{j,t} \sim N(0, \sigma_j)$, where σ_j has a gamma distribution fitted to CAPM residual variances
4. Market updates its beliefs about each fund's alpha:

$$\begin{aligned} m_{j,t+1} &= (1 - w_{j,t})m_{j,t} + w_{j,t} r^e_{j,t+1}, \\ v_{j,t+1} &= \sigma_j^2 w_{j,t}, \end{aligned}$$

where $w_{j,t} = \frac{v_{j,t}}{v_{j,t}+\sigma_j^2}$ is the weight Bayes' rule places on month $t+1$'s risk-adjusted return

Simulation evidence: Assumptions (II)

Two alternative rules for survival

1. A fund disappears if $m_{j,t}$ falls below -0.55%
2. A fund disappears if $\Pr(\alpha_j > 1.3\% \mid m_{j,t}, v_{j,t}) < 0.05$

Parameters chosen to match the 10-year survival rate

- Run randomized-exit simulation next to actual simulation
- All random draws (e.g., fund alphas, idiosyncratic shocks, ...) the same across these simulations
- Except: when a fund disappears from the actual simulation, a *randomly* chosen fund disappears from the RE simulation

  → Exactly match the amount of "luck" in the simulations

Simulation evidence: Results

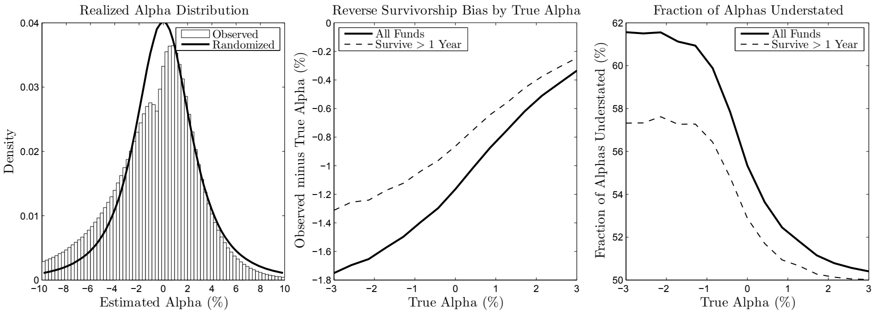

It is not just about the mean

- The salient effect is the shift in the distribution's mean
  - But this could be circumvented (under some assumptions)
- The general problem is the distortion of the entire distribution
  - We take mass from somewhere in the true distribution...
    ...and move it to the observed distribution's left tail
  - Because idiosyncratic shocks very volatile relative to the variance of true alphas, $\mathrm{var}(\tilde{\varepsilon}_{j,t}) \gg \mathrm{var}(\alpha_j)$, the distortion extends over the entire distribution
- To undo this distortion, have to model fund survival as a function of past returns

Structural Estimation

Structural model: Assumptions (I)

1. Alphas drawn from a true distribution $N(\mu, \sigma)$ that is known to the market
2. Each fund generates returns that the market risk adjusts:

$$\tilde{r}^e_{j,t} \equiv \alpha_j + \tilde{\varepsilon}_{j,t},$$

where $\tilde{\varepsilon}_{j,t}$ is normal with mean zero and variance σ_e^2

$\rightarrow$ σ_e^2 represents two sources of uncertainty: idiosyncratic shocks and AP model estimation uncertainty

$\rightarrow$ The AP model is correct: $\mathrm{E}[\tilde{r}^e_{j,t}] = \alpha_j$

Structural model: Assumptions (II)

3. Market uses the Bayes' rule to update beliefs about α_js:

$$\begin{aligned} m_{j,t+1} &= (1 - w_{j,t}) m_{j,t} + w_{j,t} r^e_{j,t+1} \\ v_{j,t+1} &= \sigma_j^2 w_{j,t} \end{aligned}$$

4. Fund exit probability specified exogenously
   - Market computes the amount of mass, $p_{j,t}$, in the posterior distribution below some level $\bar{\alpha}$
   - The hazard rate for fund j in month t is

$$H(p_{j,t},\, t) = \Pr\left(\tilde{u}_{j,t} \leq \gamma_0 + \gamma_1\, p_{j,t} + \gamma_2\, p_{j,t}^2 + \gamma_3\, t\right),$$

   where $\tilde{u}_{j,t}$s are independent draws from a uniform distribution
   - Fixed, time-independent threshold a special case
   - Estimate $\bar{\alpha}$, γ_0, γ_1, γ_2, and γ_3 from the data

Estimation (Calibration)

- Eight parameters to estimate: μ, σ^2, σ_e^2, and the five hazard-rate parameters
- Match the following moments between the model and data:
  1. Average alphas for surviving funds in years 3–4, ..., 9–10
  2. Average alphas for dead funds in these same intervals
  3. Fraction of funds disappearing in these same intervals
  4. Mean and variance of the observed alpha distribution
- These moments instruct the model to
  - Match the levels of alphas for surviving and perishing funds
  - Match the changes in these alphas as funds mature
  - Mutual funds in the model drop out at the right rate
  - Push the model's observed alpha distribution close to the data

Moments in data vs. model

Moment Condition	CAPM Data	CAPM Model	Four-Factor Data	Four-Factor Model
Alpha Estimates for Disappearing Funds				
Years 3–4	−4.83	−5.04	−4.12	−4.56
Years 5–6	−3.58	−3.40	−3.37	−3.03
Years 7–8	−2.55	−2.41	−2.40	−2.13
Years 9–10	−1.88	−1.72	−1.88	−1.50
Alpha Estimates for Surviving Funds				
Years 3–4	−0.11	0.27	0.12	−0.13
Years 5–6	0.35	0.45	0.17	0.02
Years 7–8	0.73	0.60	0.15	0.13
Years 9–10	0.75	0.70	0.03	0.20
Fraction of Funds Disappearing				
Years 3–4	6.76	6.30	6.77	6.73
Years 5–6	8.14	8.19	8.14	8.29
Years 7–8	8.54	8.85	8.54	8.49
Years 9–10	9.06	9.02	9.06	8.37
Observed Alpha Distribution				
Mean	−0.55	−0.59	−0.92	−0.91
Standard Deviation	4.83	5.00	4.03	4.01
SD of Idiosyncratic Risk	9.36	9.34	7.13	7.14

Structural parameter estimates

Parameter	CAPM		Four-Factor	
	EST	SE	EST	SE
μ	−0.129	(0.079)	−0.494	(0.092)
σ	2.756	(0.452)	2.112	(0.375)
$\bar{\alpha}$	−2.846	(3.931)	−2.098	(2.195)
σ_e	9.340	(0.125)	7.142	(0.095)
γ_0	0.286	(1.044)	0.767	(1.968)
γ_1	−0.035	(0.102)	−0.073	(0.126)
γ_2	0.001	(0.000)	0.002	(0.001)
γ_3	−0.043	(0.120)	−0.119	(0.249)
J-test, χ^2 (p-value)	9.78	(0.202)	12.92	(0.074)

Observed vs. true distributions in the model and data

Model	Mean	Percentiles				
		5%	25%	50%	75%	95%
CAPM						
Data	−0.547	−8.568	−2.567	−0.242	1.908	5.969
Model						
Observed	−0.587	−9.293	−3.261	−0.188	2.615	6.662
True	−0.129	−4.662	−1.988	−0.129	1.730	4.404
Four-Factor Model						
Data	−0.920	−7.665	−2.848	−0.767	0.956	5.202
Model						
Observed	−0.910	−7.988	−2.921	−0.539	1.606	4.746
True	−0.494	−3.968	−1.919	−0.494	0.930	2.980

Comments on estimation results

- The mean bias ~50 basis points per year
  - Significant distortion in the shape of the alpha distribution
  - True distribution symmetric around -0.13% but the observed distribution's 5^{th} and 95^{th} percentiles -9.3% and 6.7%!
- Estimates of prevalence of skill?
  - 12% of funds have four-factor model alphas greater than 2% per year
- The low dispersion of the alpha distribution (2.1%) relative to idiosyncratic volatility (7.1%) is a problem for investors:
  - Example: It takes 11 years of data for the variance of the posterior distribution to decrease by half!
  - *Very difficult* to identify good funds from past returns alone

Remedies

1. Discard funds that survive for fewer than m' months
   - What is the correct choice of m'?
   - Does not correct the shape of the distribution
2. Estimate alphas using only one year of data for each fund
   - Very noisy estimates, low power
3. Evaluate funds not by $\hat{\alpha}$ but by $t(\hat{\alpha})$
   - Reduces the bias by increasing the denominator in $t(\hat{\alpha}) \equiv \frac{\hat{\alpha}}{\mathrm{SE}}$
   - But it cannot offset the bias
   - This procedure gives more weight to funds that survive longer
4. Estimate alphas using the average alpha of fund holdings
   - Quarterly data; misses trading etc. costs
   - Window dressing?

Conclusions

- Direct survivorship bias:

  Because mutual funds often disappear following poor performance, conditioning on survival overstates performance

- But performance can be poor not only because the true alpha is low but also because of negative idiosyncratic shocks

  $\Rightarrow$ Reverse survivorship bias

- Start without dead funds and then begin adding them back
  - Direct survivorship bias goes down
  - Reverse survivorship bias goes up
  - With all dead funds in, the mean bias about 60 bps per year
- Shifts mass towards the observed alpha distribution's left tail
  - Entire distribution distorted because $\mathrm{var}(\tilde{\varepsilon}_{j,t}) \gg \mathrm{var}(\alpha_j)$

Households' and pension funds' alphas, CEOs' abilities... and two-stage procedures?

- Run a panel regression with CEO fixed effects to extract measures of CEOs' abilities
- Explore determinants of CEO skill in a 2^{nd}-stage regression
- What if there is cross-sectional variation in reverse survivorship bias?
  - Examples: differences in how sensitive CEOs are to bad performance, differences in idiosyncratic volatility, ...
- The 2^{nd}-stage estimates will pick up this variation
  - Suppose some variable correlates with how sensitive CEO "survival" is to bad performance (e.g., tenure?)
  - It will appear to explain variation in CEO skill
  - Errors-in-variables problem: the noise in the dependent variable is correlated with the RHS variables in the second stage

Other implications

- The reverse survivorship bias increases in
  1. Return volatility
  2. Dispersion in survival times
  3. Correlation between returns and survival
- Barras et al. (2010): the proportion of skilled managers decreases markedly between 1996 and 2006
- If the mutual fund landscape grew more competitive, the correlation between returns and survival might have increased
- If idiosyncratic volatility increased, the bias would have been amplified

  → The deterioration in alphas may have been partly due to a strengthening reverse survivorship bias