Confidence of the Trembling Hand: Bayesian Learning with Data Poor Stocks

Jorge Holzer and Qian Qu

La Paz, Mexico

March 23, 2017

Motivation

Data-Poor Stocks

- Increasing interest in the development of tools to manage data-poor stocks.
- Managers must often make decisions irrespective of data availability or completeness of scientific understanding.
- MSA requires annual catch limits (ACLs) to prevent overfishing for most federally managed species, including data-poor stocks.
- Similar requirements exist in Australia and New Zealand.
- Catch-based approaches are typically used when there are insufficient data for estimating overfishing limits (OFL) with more sophisticated methods (i.e. stock assessment models).

- These approaches set the OFL as the average (or median) catch over a reference period, and calculate the ACL according to downward adjustment based on uncertainty about stock status.

- These methods are currently used in a number of fisheries: Greenland halibut, Snowy grouper, Atlantic mackerel, Red crab, Golden king crab, Flathead sole, etc (Newman et al. 2015).

- We compare these methods ("steady" hand) with an alternative strategy: perturbations in the form of small, temporary and intermittent increases to ACLs ("trembling" hand).

The Resource:

- Stock grows according to a logistic model up to a random disturbance

$$G_t = r_0 S_{t-1}\left(1 - \frac{S_{t-1}}{\Omega}\right) + \eta_t$$

  where r_0 denotes the intrinsic growth rate, Ω the carrying capacity and $\eta_t \sim iid\ N(0, \sigma_\eta^2)$.

- The evolution of the stock is governed by $S_t = S_{t-1} + G_t - h_t$.

- At each date a noisy signal $y_t = S_t + \epsilon_t$ is generated, where $\epsilon_t \sim iid\ N(0, \sigma_\epsilon^2)$.

The Regulator:

- Knows all parameters but r_0. Regulator's beliefs at $t-1$ given by the distribution function Q_{t-1}.
- Manages the resource with a fishery-wide ACL ($\bar{H}_t$).
- $\bar{H}$ chosen so that the maximum sustainable yield is only exceeded with probability α

$$Pr[MSY_{t-1} \leq \bar{H}_t] = \alpha \;\Leftrightarrow\; \bar{H}_t = Q_{t-1}^{-1}(\alpha)\frac{1}{4}\Omega$$

- Manager updates beliefs every $\mathcal{T}$ periods (i.e. using signal series of length $\mathcal{T}$).

The Industry:

- Industry profits are given by $\pi_t = p_t h_t - c(S_t, h_t)$, where $c_S < 0,\ c_h > 0$.
- Industry sets $h_t = \min[\bar{H}, H_0]$, where $\bar{H}$ is the annual catch limit set by the regulator and H_0 the zero-profit aggregate catch level.

Intuition: Two Types
Additional Channel of Information

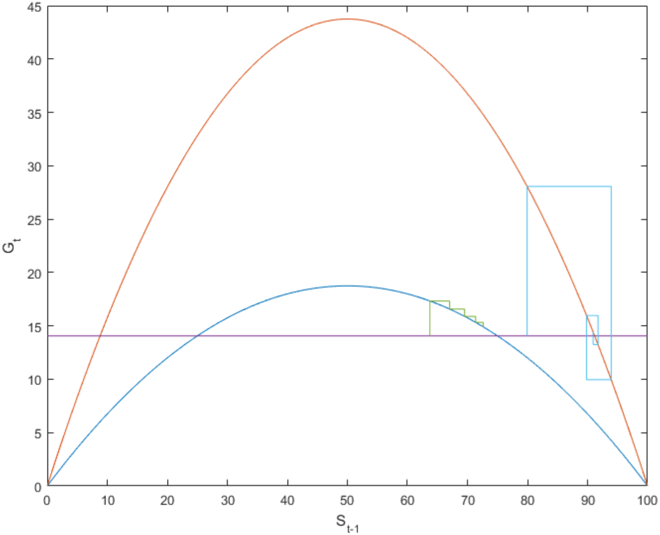

Intuition: Two Types

Type A Difference

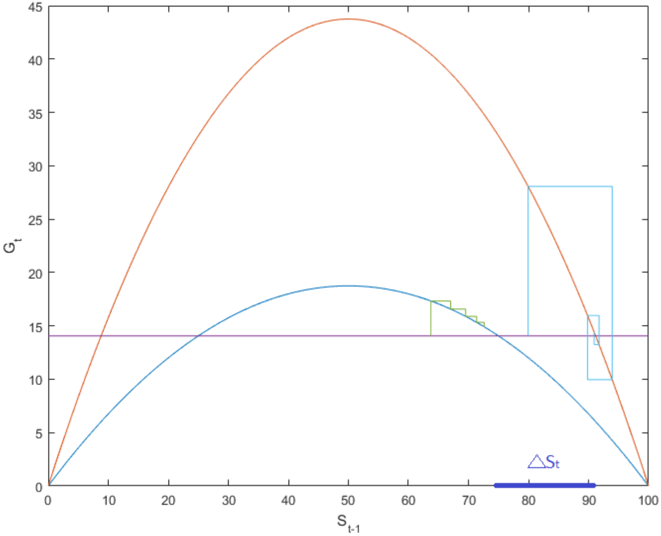

Intuition: Two Types

Type B Difference

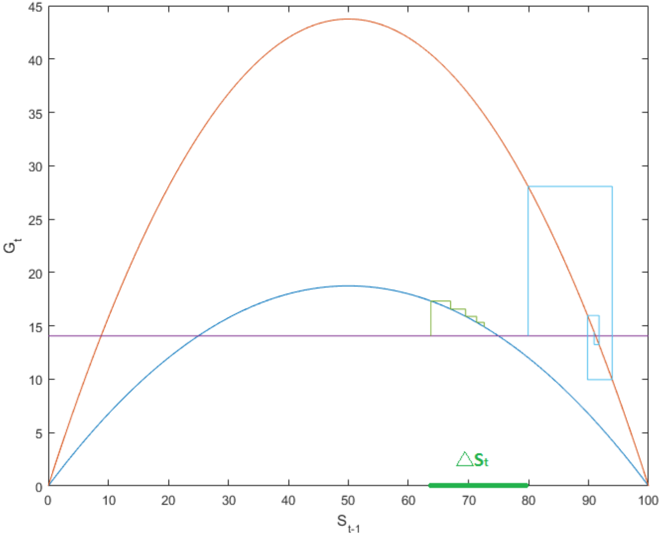

Intuition: Two Types

Type B Difference

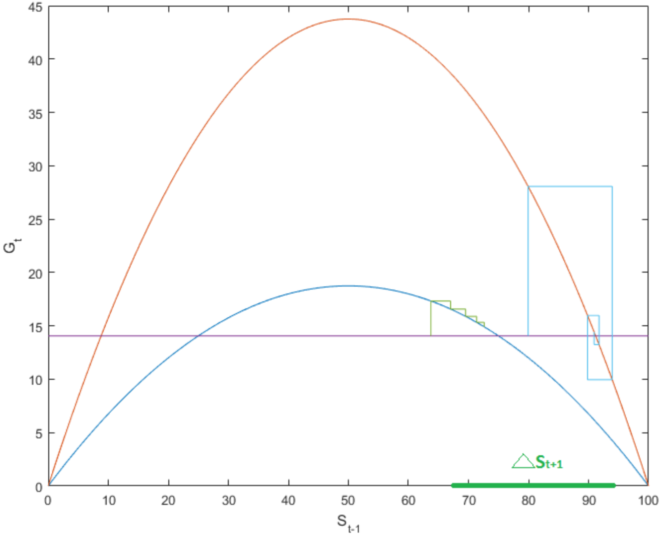

Intuition: Two Types

Type B Difference

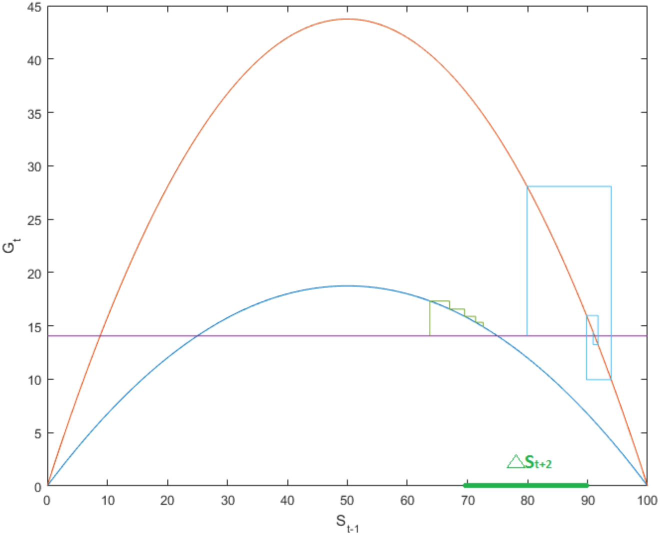

Intuition: Two Types

Type B Difference

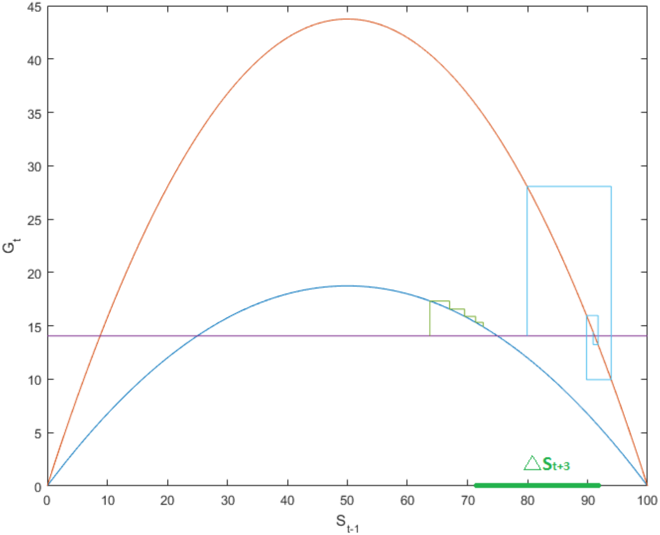

Intuition: Two Types

Type B Difference

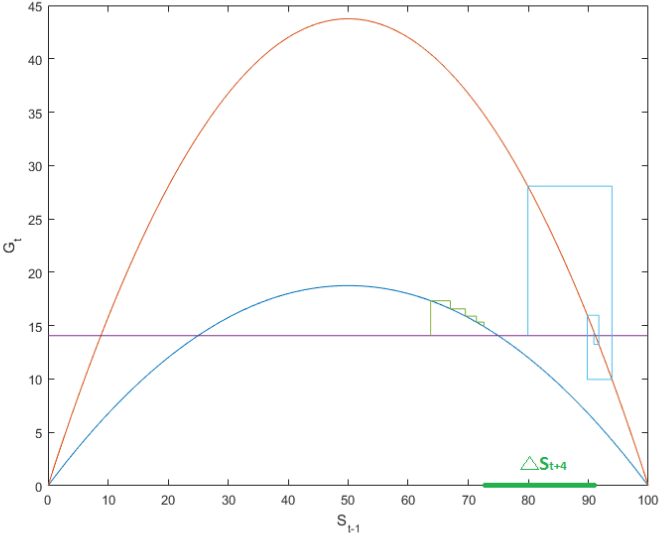

- Two convergence patterns to the high steady state after perturbation: r^L approaches from one side while r^H fluctuates with diminishing amplitude.
- The intrinsic growth rate parameter governs both the location of the high steady state and its surrounding dynamics.
- Steady hand: relies solely on the difference in high steady state location to distinguish r^H and r^L (Type A difference).
- Trembling hand: unseals the difference in dynamics (Type B difference).

Intuition: Two Types

Case I

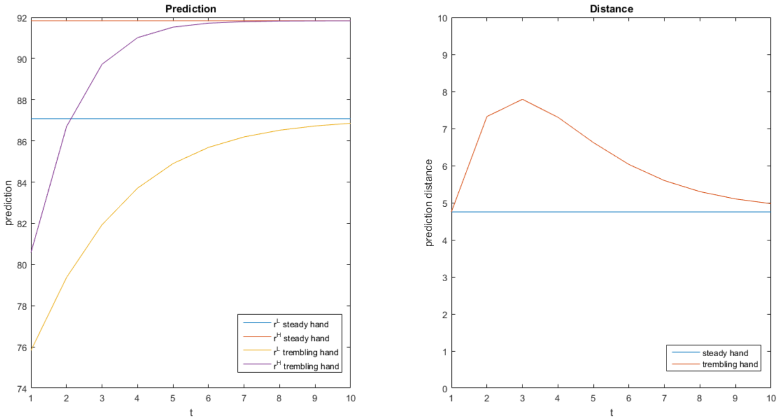

- Convergence from one side of the steady state: r^H generates faster convergence path than r^L
- Trembling always has a bigger gap than steady.
- The superiority of trembling diminishes over time.
- In the actual work, cumulative confidence difference for all $\mathcal{T}$ days is computed by weighting the difference of each day by a measure of accuracy and summing over the horizon.
- Cumulative confidence difference increases over time at a decreasing rate.
- Type A difference always adds to type B difference.

Intuition: Two Types

Case II

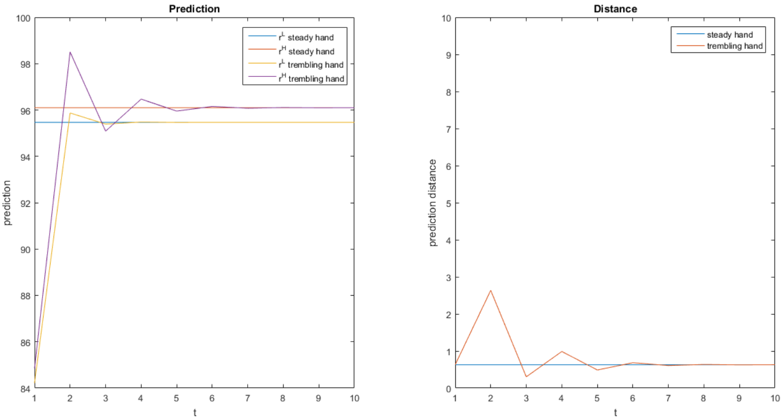

Intuition: Two Types

Case II

- Converge from two sides: r^H at a slower rate than r^L.
- Under the trembling hand, bigger and smaller gaps alternate.
- Confidence gain and loss alternate.
- Cumulative confidence difference does not grow monotonically.
- Cumulative confidence is (almost) always positive: (1) alternation begins with information gain; (2) sizes of gain/loss diminish over time.
- Type A and type B difference add to each other at one period, work against each other at the next.
- When type A and type B difference are antagonistic, type A difference dominates.

Intuition: Two Types

Case III

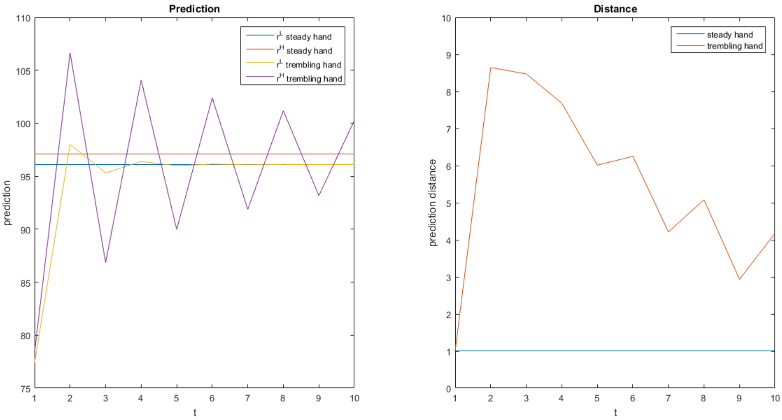

Intuition: Two Types

Case III

- Convergence from two sides: r^H at a lower rate than r^L.
- Confidence gain every period. The size of gain diminishes over time. Cumulative information difference monotonically increases.
- Type A and type B differences oscillate from being synergistic to being antagonistic (i.e. there is internal friction caused by the oscillating pattern of convergence).
- When two types of differences work against each other, type B difference overwhelms type A.

General Case

A Continuum of Types

Setting

- Candidate parameter values on a continuum
- Consider the two management strategies:
  - (i) $h_t = \bar{H}, \ \forall t$
  - (ii) $h_t = \bar{H} + \Delta_t$ where at least one $\Delta_t \neq 0$ for $1 \leq t \leq \mathcal{T} - 1$

Approach

- Linearize around the high steady state to obtain the state space representation
- Apply the Kalman filter to construct the likelihood
- Analytically derive the difference in fisher information for the two cases
- Explore the range of information gain

State-Space Representation

Linearize around the high steady state and solve for the steady state representation:

$$
\begin{aligned}
y_t =& A(r) + \tilde{S}_t + \epsilon_t \\
\tilde{S}_t =& B(r)\,\tilde{S}_{t-1} - \Delta_t + \eta_t
\end{aligned}
$$

where

$$
\begin{aligned}
A(r) = S^*_{high}\left(r,\ \Omega,\ \bar{H}\right) =& \frac{r + \sqrt{r^2 - 4\dfrac{r}{\Omega}\bar{H}}}{2r/\Omega} \\
B(r) =& r - \frac{2rA(r)}{\Omega} + 1
\end{aligned}
$$

- Let $Y_t = \{y_s\}_{s=1}^{t}$ denote signals up to date t, $a_t = E[\tilde{S}_t | Y_{t-1}]$ the one-step predicative mean of stock deviation from the steady state, and $p_t = var[\tilde{S}_t | Y_{t-1}]$ the predicative variance.

- Applying the Kalman algorithm, a_t and p_t evolve as follows:

$$
\begin{aligned}
a_t =& B(r)\, a_{t-1} + B(r)\, K_{t-1} e_{t-1} - \Delta_t \\
p_t =& B^2(r)\, p_{t-1} - B^2(r)\, K_{t-1} p_{t-1} + \sigma_\eta^2 \\
e_t =& y_t - A(r) - a_t \\
F_t =& var(e_t) = p_t + \sigma_\epsilon^2 \\
K_t =& \frac{p_t}{F_t}
\end{aligned}
$$

Likelihood and Fisher Information

- The log-likelihood takes the form:

$$\ln \mathscr{L}\left(Y_{\mathcal{T}} \mid r, \{\Delta_t\}\right) = -\frac{\mathcal{T}}{2}\ln 2\pi - \frac{1}{2}\sum^{\mathcal{T}} \ln F_t - \frac{1}{2}\sum^{\mathcal{T}} e_t^2 F_t^{-1}$$

- We use Fisher information $\mathcal{I}_{\mathcal{T}}(r) = -E\left[\frac{\partial^2}{\partial r^2}\ln \mathscr{L}\left(Y_{\mathcal{T}} \mid r, \{\Delta_t\}\right)\right]$ to measure the confidence the regulator can have in its estimate (i.e. information the linearized system accumulates about r up to date $\mathcal{T}$).

- $\Delta\mathcal{I}_{\mathcal{T}}(r)$ between the two management strategies is equivalent to the difference in $\sum^{\mathcal{T}} F_t^{-1} E\left[A'(r) + \frac{\partial}{\partial r} a_t\right]^2$, a measure of distance between true and neighboring false predictions.

Fisher Information Difference

Exogenous Variation

The measure of prediction distance can be written as the difference in exogenous and deterministic variation:

$$
\begin{aligned}
\Delta\mathcal{I}_{\mathcal{T}}(r) = & \sum_{t=1}^{\mathcal{T}} F_t^{-1} \left[\left(\sum_{m=2}^{t} C_{m,t} \left(-\Delta_m\right) + D_t a_1^{TR} \right) B'(r) + A'(r)\left(1 - M_t\right) \right]^2 \\
& - \sum_{t=1}^{\mathcal{T}} F_t^{-1} \left[D_t a_1^{ST} B'(r) + A'(r)\left(1 - M_t\right) \right]^2
\end{aligned}
$$

where

$$
\begin{aligned}
C_{m,t} &= \left(\frac{\partial B^{t-m}}{\partial B} \right) + \sum_{s=m}^{t-1} \left(-B^{t-s} K_s\right) C_{m,s}, \quad C_{m,m} = 0, \quad t \geq m+1 \\
D_t &= \left(\frac{\partial B^{t-1}}{\partial B} \right) + \sum_{s=1}^{t-1} \left(-B^{t-s} K_s\right) D_s, \quad D_1 = 0, \quad t \geq 2 \\
M_t &= \left(B - BK_{t-1}\right) M_{t-1} + BK_{t-1}, \quad M_1 = 0, \quad t \geq 2
\end{aligned}
$$

Posterior Beliefs

- To obtain the posterior, use the fact that by Bayes' rule the posterior is proportional to the likelihood multiplied by the prior:

$$\ln \mathscr{L}\left(r|Y_{\mathcal{T}}\right) = \ln \mathscr{L}\left(Y_{\mathcal{T}}|r,\{\Delta_t\}\right) + \ln\left(\frac{d}{dr}Q_0(r)\right) \tag{1}$$

- Maximization of (1) results in the maximum likelihood estimate $\hat{r}_{\mathcal{T}}$, and the corresponding inverse of the information matrix evaluated at the MLE estimate, $\hat{\sigma}^2_{\mathcal{T}} = -\left(\frac{\partial^2}{\partial r^2} ln\, \mathscr{L}(\hat{r}_{\mathcal{T}})\right)^{-1}$.

- Thus, the manager's beliefs on the intrinsic growth parameter at time $\mathcal{T}$ are given by $R_{\mathcal{T}} \sim N(\hat{r}_{\mathcal{T}}, \hat{\sigma}^2_{\mathcal{T}})$.

Case I: confidence gain every period

- $\sum_{m=1}^{t} C_{m,t}\left(-\Delta_m\right) B'\left(r\right)$ and $A'\left(r\right)\left(1-M_t\right)$ are of the same sign
- Falls into case I if $B'(r) > 0 \Leftrightarrow \bar{H} > \dfrac{r^2-1}{r^2}\left(\dfrac{1}{4}r\Omega\right)$

Case II: confidence gain and loss alternate

$\sum_{m=1}^{t} C_{m,t}\left(-\Delta_m\right) B'\left(r\right)$ and $A'\left(r\right)\left(1-M_t\right)$ are of different signs and $\left|\sum_{m=1}^{t} C_{m,t}\left(-\Delta_m\right) B'\left(r\right)\right| < 2\left|A'\left(r\right)\left(1-M_t\right)\right|$

Case III: confidence gain every period

$\sum_{m=1}^{t} C_{m,t}\left(-\Delta_m\right) B'\left(r\right)$ and $A'\left(r\right)\left(1-M_t\right)$ are of different signs and $\left|\sum_{m=1}^{t} C_{m,t}\left(-\Delta_m\right) B'\left(r\right)\right| > 2\left|A'\left(r\right)\left(1-M_t\right)\right|$

A Numerical Example ($\Delta = \sigma_\eta$)

Fisher Information Gain

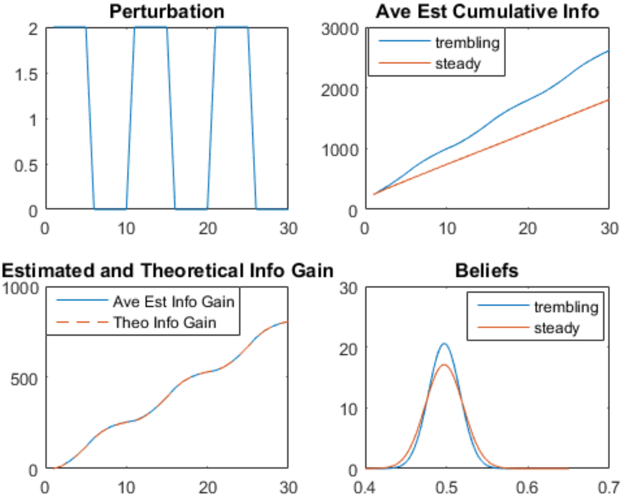

Stock NPV Expectation

Different perturbation sizes and safety valves

- Safety valve: if $\dfrac{y_t}{y_{t-1}} < \phi$, remaining perturbations are canceled.

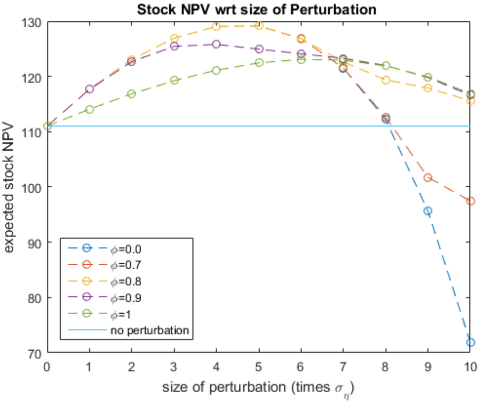

Stock NPV Expectation

Different perturbation sizes and safety valves

- Inverted U-shape wrt size of perturbation. As the size of perturbation increases, expected stock NPV first increases then decreases.
- Safety valve too loose ($\phi \leq 0.7$), under large perturbation trembling is inferior to steady.
- Safety valve too tight ($\phi \geq 0.9$), ensures that trembling always does better than steady at the cost of reduced profitability when perturbation is small.
- When the safety valve is set just right ($\phi = 0.8$), trembling with the valve $\succeq$ max{trembling without the valve, steady hand} for all sizes of perturbation.

Stock NPV Distribution

Different perturbation sizes and safety valves

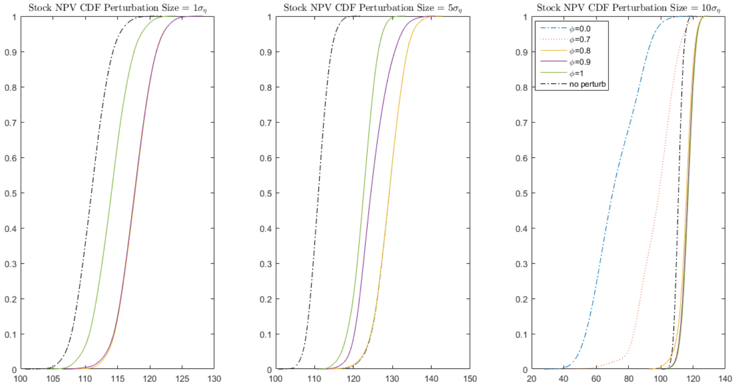

Stock NPV Distribution

Different perturbation sizes and safety valves

- Left and middle panel: tight valves unnecessarily reduces profitability. Perturbation without the valve first order stochastically dominates all others.
- Right: safety valve necessary to ensure that trembling does no worse than steady. Need $\phi \geq 0.9$ for trembling to first order stochastically dominate steady.

Trembling-Updating Equivalence

Assessment Cost Equivalent to 10% Annual Revenues under Steady Hand

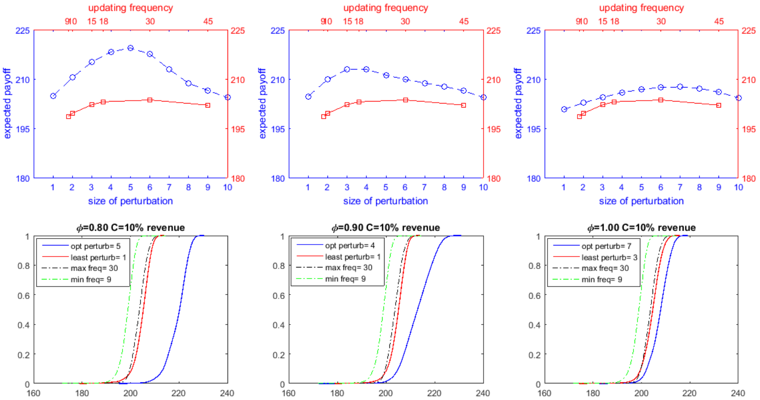

Trembling-Updating Equivalence

Assessment Cost=$0

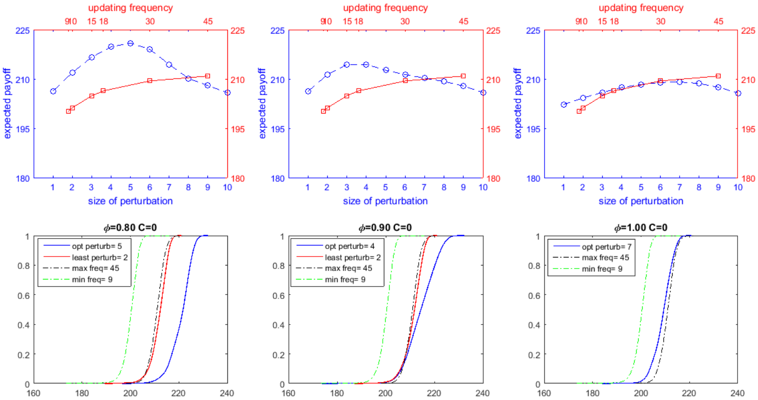

Cost Equivalence

- Upper Panel: expected payoff wrt various sizes of perturbation and updating frequencies.
- Lower Panel: (1) perturbation size that yields the max payoff; (2) smallest perturbation that dominates all updating frequencies; (3) updating frequency that yields the maximum payoff; (4) updating frequency that yields the minimum payoff.
- In most cases relatively small perturbation dominates all updating frequencies.
- Only when updating is costless and valve is extremely stringent, the optimal payoff of perturbation is inferior to that of updating.

Final remarks

- Small, temporary and intermittent increases to ACLs between stock assessments (combined with a safety valve) can increase the permanent value of the stock.
- These perturbations translate into a more accurate estimation model in the next assessment, which leads to ACLs closer to the true MSY thereafter.
- On the other hand, increases in ACLs imply a higher management risk in the form of overfishing.
- The net effect will depend on the interaction between type A and type B differences.