Adaptive/Learning Quantum Control via Extremum Seeking Feedback

Robert Kosut

SC Solutions, Sunnyvale, CA

Presented at
BIRS Workshop on Control of Open Quantum Systems
Banff April 3-8, 2011

Research supported by DARPA Grant Award FA9550-09-1-0710

On Learning

SEEKER “How do I gain **Good Judgement**?”

On Learning

SEEKER “How do I gain **Good Judgement**?”

SAGE “You must first acquire **Wisdom**.”

On Learning

SEEKER “How do I gain **Good Judgement**?”

SAGE “You must first acquire **Wisdom**.”

SEEKER “How do I gain **Wisdom**?”

On Learning

SEEKER	"How do I gain **Good Judgement**?"
SAGE	"You must first acquire **Wisdom**."
SEEKER	"How do I gain **Wisdom**?"
SAGE	"From **Bad Judgement**."

On Learning

SEEKER	"How do I gain **Good Judgement**?"
SAGE	"You must first acquire **Wisdom**."
SEEKER	"How do I gain **Wisdom**?"
SAGE	"From **Bad Judgement**."

$$\Downarrow$$

First rule of learning: **Make Mistakes**

On Learning

SEEKER	“How do I gain **Good Judgement**?”
SAGE	“You must first acquire **Wisdom**.”
SEEKER	“How do I gain **Wisdom**?”
SAGE	“From **Bad Judgement**.”

$$\Downarrow$$

First rule of learning: **Make Mistakes**

Second rule: FEEDBACK

Adaptive/Learning Feedback: Open-Loop

Quantum system in the k-th iteration:

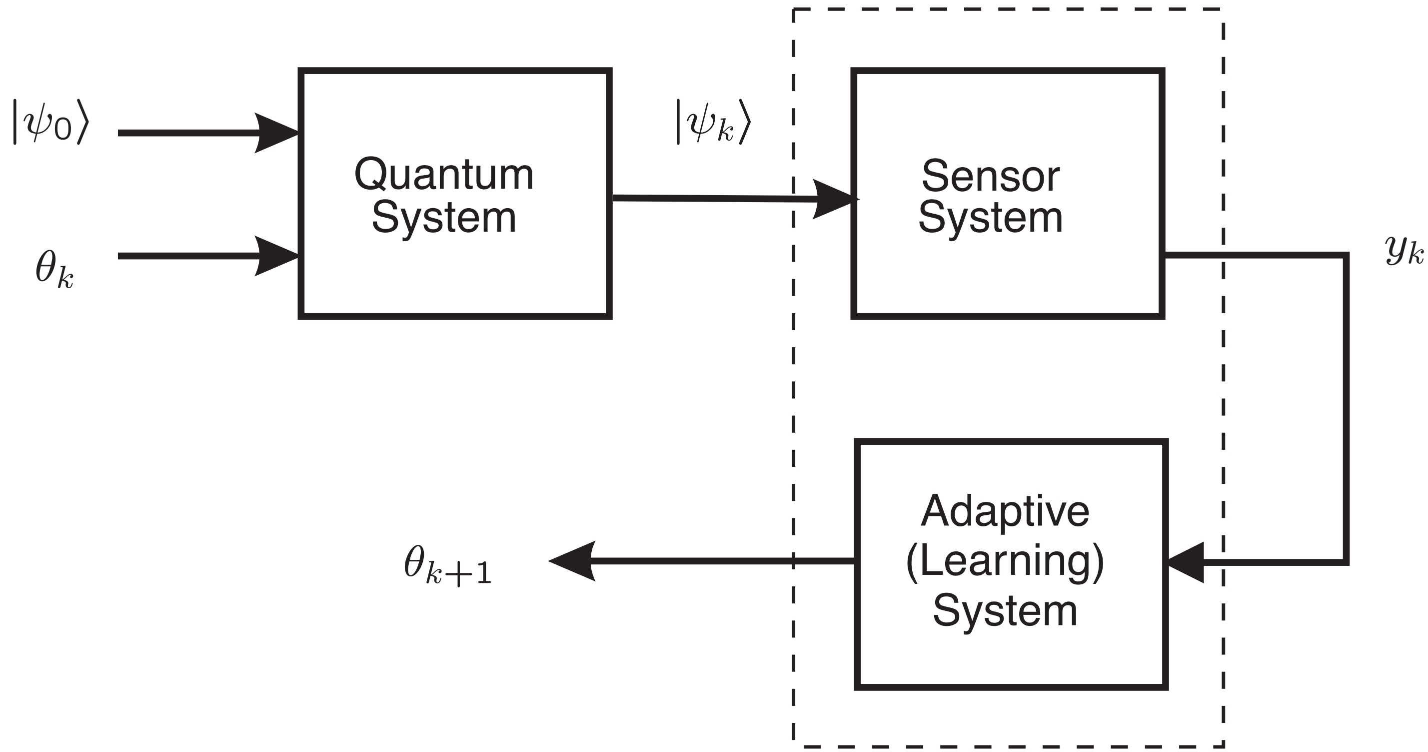

- initial state $|\psi_0\rangle$ is identical in each iteration, *i.e.*, system is "re-freshed" at every iteration.
- $\theta_k \in \mathbf{R}^n$ are parameters that define k-th control.
- sensor and feedback systems inside the dashed box are "classical".
- sensor system is measurement probe/meter together with data processing capabilities.
- **feedback (adaptive/learning) system generates the control θ_{k+1} for the next iteration based on past measurements and controls.**
  - how simple and comprehensible can this be, and will that help?

Example system

- Hamiltonian and associated unitary,

$$H(\theta) = \theta Z + \delta X = \begin{bmatrix} \theta & \delta \\ \delta & -\theta \end{bmatrix}, \; U(\theta) = \exp\{-iH(\theta)\}$$

  - control θ in Pauli-Z with a perturbation δ in Pauli-X.
  - the two eigenvalues of $H(\theta)$ and $U(\theta)$ are

$$\lambda\{H(\theta)\} = \pm\sqrt{\theta^2 + \delta^2}, \; \lambda\{U(\theta)\} = \exp\left\{\pm i\sqrt{\theta^2 + \delta^2}\right\}$$

- observable is Pauli-Z

$$O = Z = \sum_{i=1}^{2} \lambda_i \Pi_i \;\Rightarrow\; \begin{cases} \lambda_1 = -1, & \Pi_1 = |1\rangle\langle 1| \\ \lambda_2 = +1, & \Pi_2 = |2\rangle\langle 2| \end{cases}$$

- initial state

$$|\psi_0\rangle = |1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

- probability outcomes

$$\begin{array}{rcl}\mathbf{Pr}\{y=-1|\theta\} & = & p_1(\theta) & = & |\langle 1|U(\theta)|1\rangle|^2 = |U_{22}(\theta)|^2 \\ \mathbf{Pr}\{y=+1|\theta\} & = & p_2(\theta) & = & |\langle 1|U(\theta)|2\rangle|^2 = |U_{21}(\theta)|^2 = 1-|U_{22}(\theta)|^2\end{array}$$

- fidelity of $U(\theta)$ with respect to identity

$$f(\theta) = |\mathbf{Tr}\, U(\theta)|^2/4 = \cos^2\sqrt{\theta^2+\delta^2}$$

Thus for any integer n, the choice of control

$$\theta_\star = \pm\sqrt{(n\pi)^2-\delta^2}$$

will *simultaneously* make $f(\theta_\star)=1$, $p_1(\theta_\star)=1$, $U(\theta_\star)\equiv$ identity.

- set $\delta=1 \Rightarrow$ fidelity maximizing control with smallest magnitude,

$$\theta_\star = \pm\sqrt{\pi^2-1} \approx \pm 2.9782 \;\Rightarrow\; U(\theta_\star) = -I_2$$

next smallest magnitude,

$$\theta_\star = \pm\sqrt{(2\pi)^2-1} \approx \pm 6.2031 \Rightarrow\; U(\theta_\star) = +I_2$$

and so on.

Conditional probability and fidelity vs. control

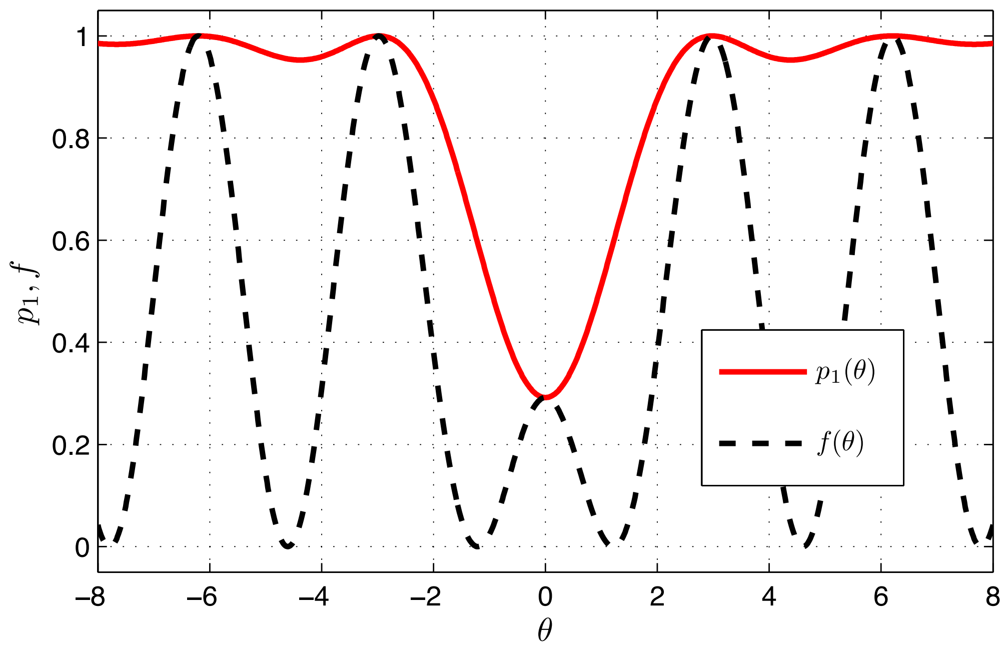

- both $p_1(\theta)$ and $f(\theta)$ vs. θ peak at the same periodic control values.
- $p_1(\theta)$ has multiple global maxima and no local maxima.
- fidelity $f(\theta)$ has both, specifically many global maxima and one local maxima.
- for quantum systems the outcomes of these functions may not instantaneously be available.

$$f(\theta_1, \theta_2) = |\mathbf{Tr}U(\theta_1, \theta_2)|^2/4, \; U(\theta_1, \theta_2) = e^{-iH(\theta_1)/2}\, e^{-iH(\theta_2)/2}$$

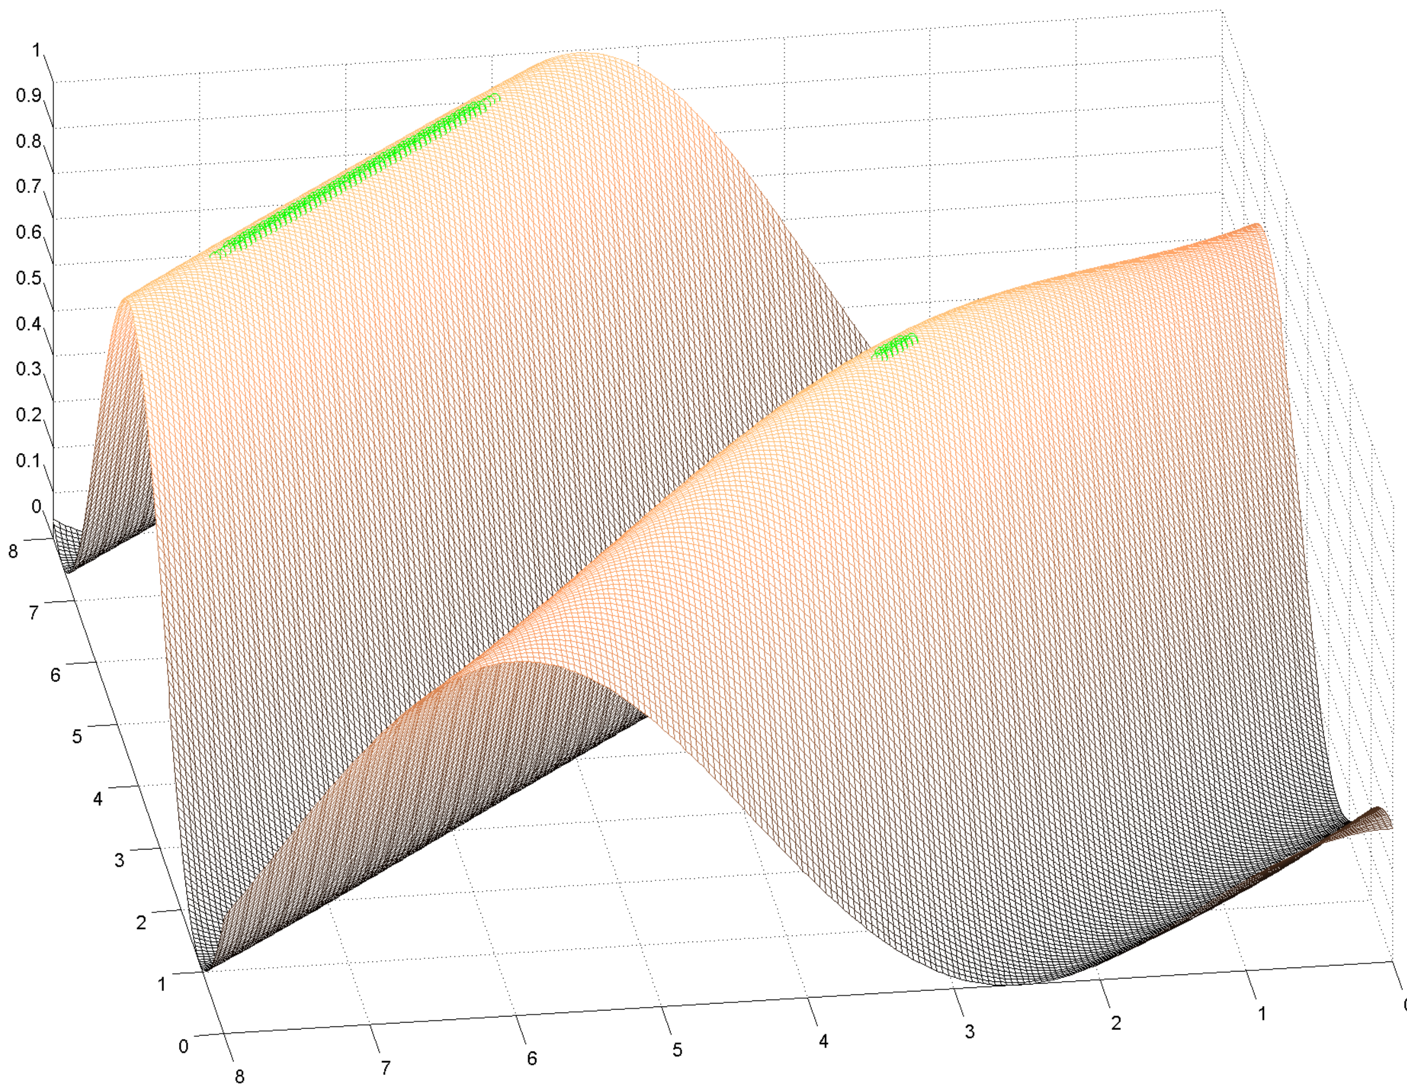

● $f(\theta_1, \theta_2) \geq 0.999$

$$f(\theta_1,\theta_2) = |\mathbf{Tr}U(\theta_1,\theta_2)|^2/4,\; U(\theta_1,\theta_2) = e^{-iH(\theta_1)/2}\, e^{-iH(\theta_2)/2}$$

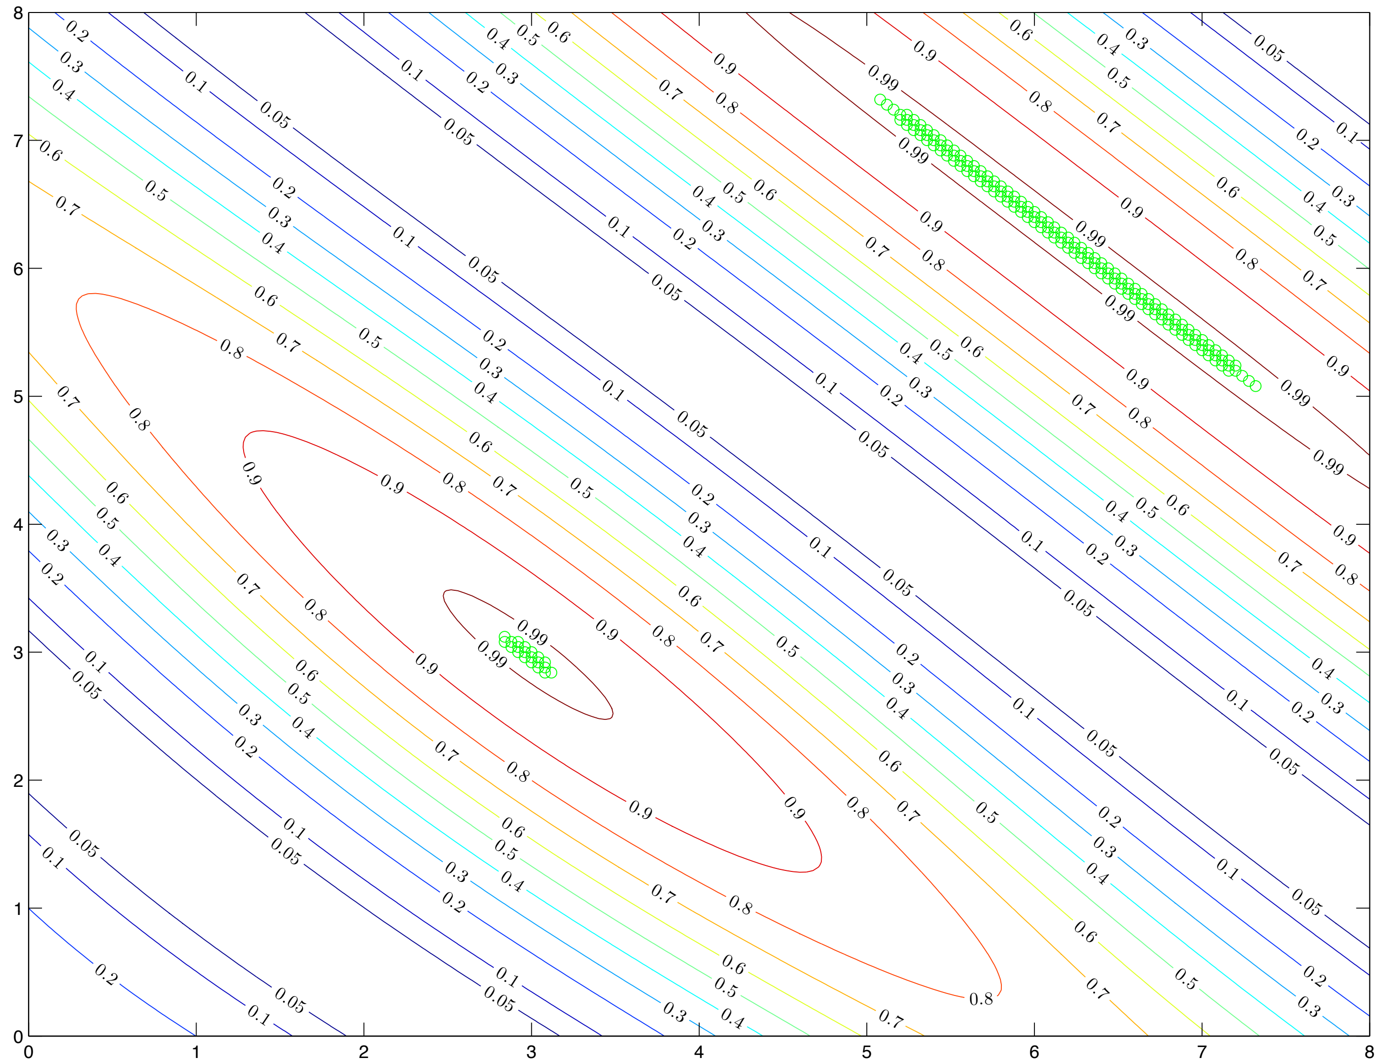

● $f(\theta_1,\theta_2) \geq 0.999$

Adaptation via discrete-time "integral-action"

- basic update: $\theta_{k+1} = \theta_k + \gamma e_k,\ \gamma > 0 \ \Rightarrow \begin{cases} \theta_k = \theta_0 + \gamma \sum_{\ell=0}^{k-1} e_\ell \\ \theta_k \to \theta_\star \ \Leftrightarrow e_k \to 0 \end{cases}$
- gradient version: $\theta_{k+1} = \theta_k - \gamma \left(\frac{\partial e_k}{\partial \theta_k}\right) e_k$ or $= \theta_k - \gamma \text{sign}\left(\frac{\partial e_k}{\partial \theta_k}\right) e_k$

- intialize the control to θ_0.
- at the start of the k-th iteration re-set the intial state to $|\psi_0\rangle$.
- with control θ_k, record the error,

$$e_k = \lambda_1 - y_k = \begin{cases} 0 & \text{with probability } p_1(\theta_k) \\ -2 & \text{with probability } 1 - p_1(\theta_k) \end{cases}$$

- update the control via,

$$\theta_{k+1} = \theta_k + \gamma e_k = \begin{cases} \theta_k & \text{with probability } p_1(\theta_k) \\ \theta_k - 2\gamma & \text{with probability } 1 - p_1(\theta_k) \end{cases}$$

With $\gamma = 0.15$, (A), (B), and (C) shows, respectively, typical responses of 1000 iterations from each of the three initial control settings $\theta_0 \in \{-1.5, -2.5, +2.0\}$.

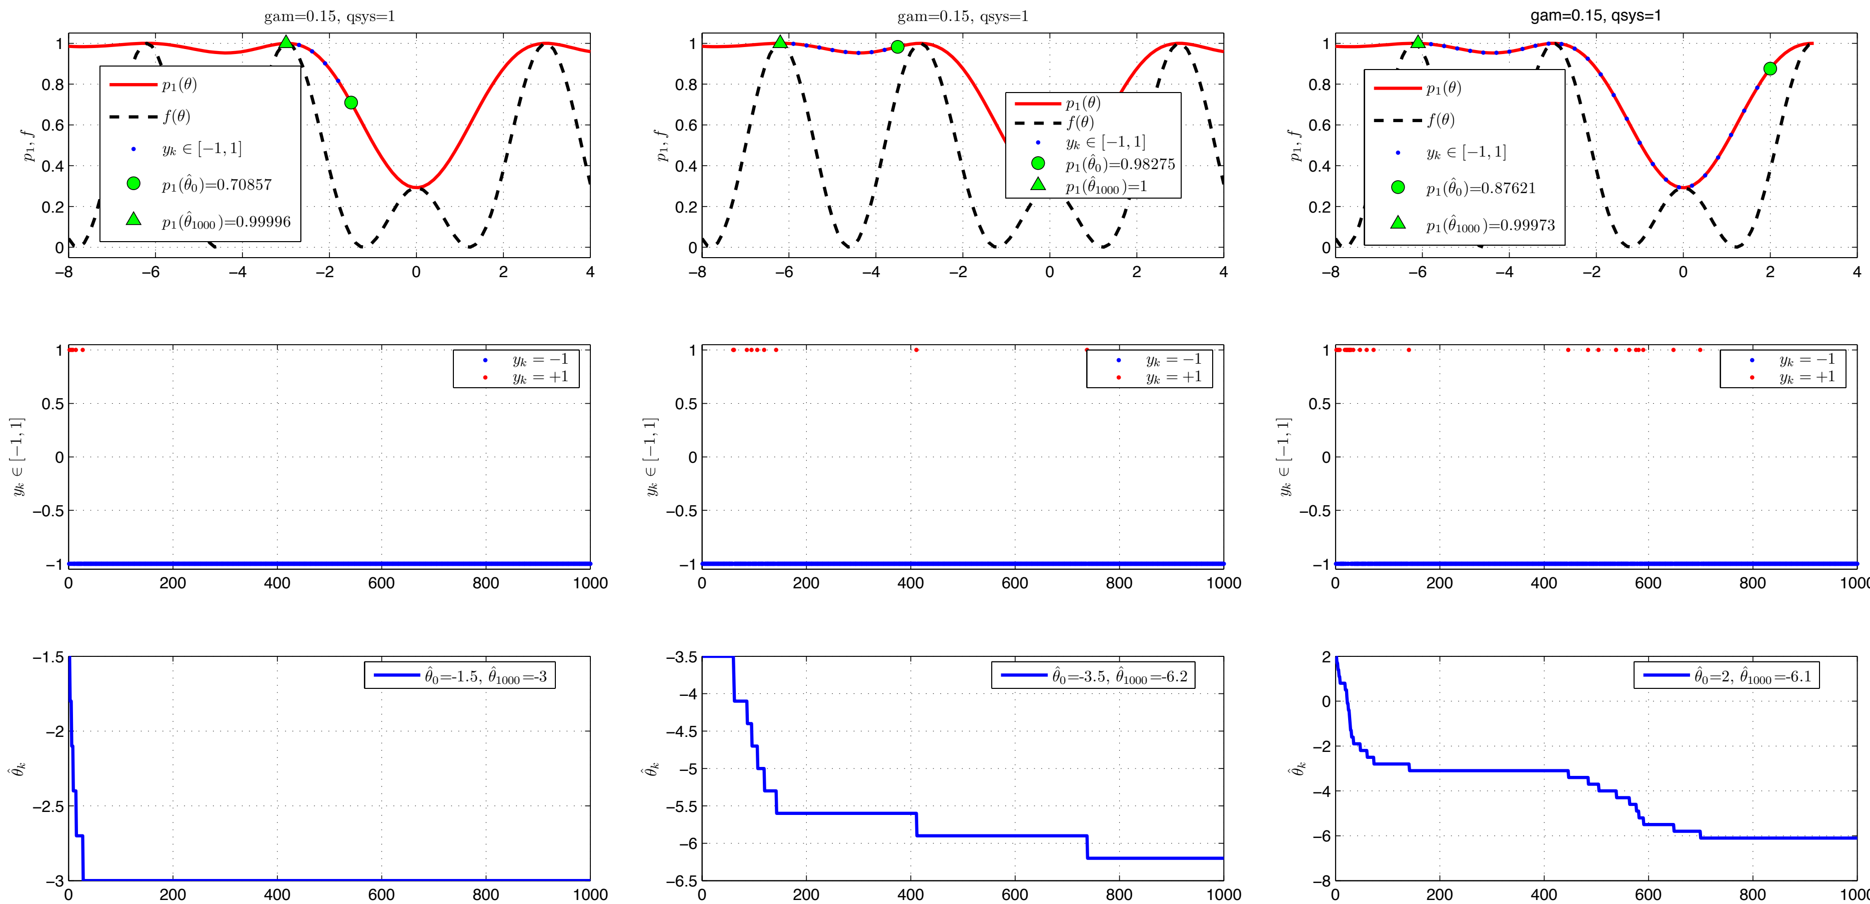

(A) $\theta_0 = -1.5$ (B) $\theta_0 = -2.5$ (C) $\theta_0 = +2.0$

Extremum Seeking Feedback*

ESF system in the k-th iteration.

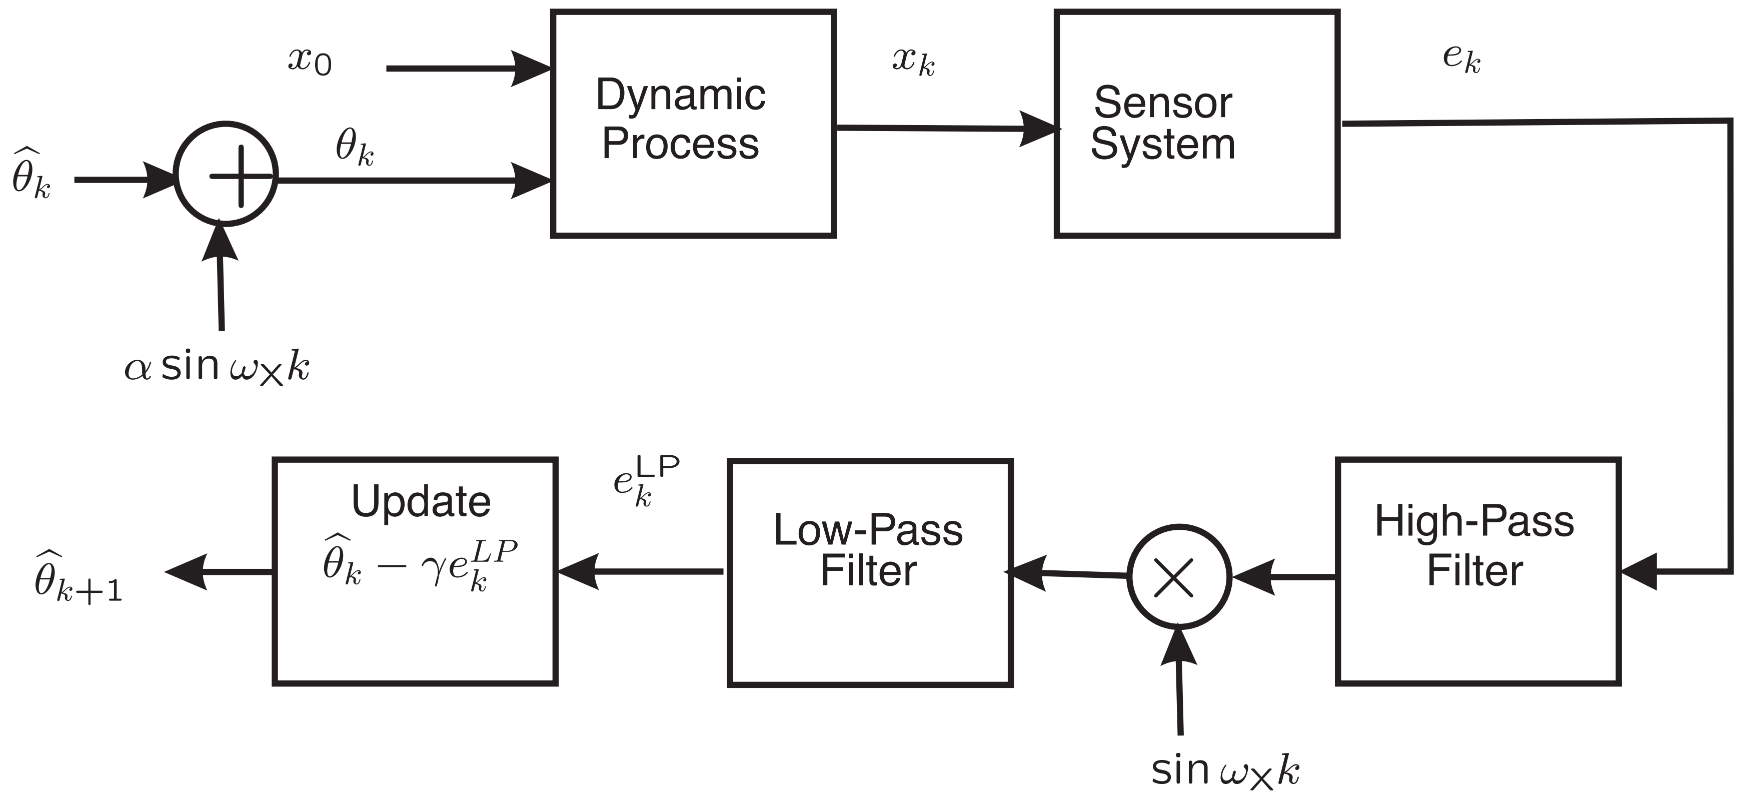

"Theory" IF:

- $e_k = e_\star + \left(\frac{\partial^2 e_k}{\partial \theta_k^2}(\theta_\star)/2\right)(\theta_k - \theta_\star)^2$ with $\frac{\partial^2 e_k}{\partial \theta_k^2}(\theta_\star) > 0$
- $\omega_{LP} \leq \omega_{HP} \leq \omega_X$
- $\gamma \alpha \frac{\partial^2 e_k}{\partial \theta_k^2}(\theta_\star) > 0$ and "small"

THEN: $\widehat{\theta}_k \to \theta_\star$

* K.B. Ariyur & M. Krstic, *Real-Time Optimization by Extremum Seeking Feedback*, Wiley, 2003.

Extremum Seeking Feedback*

ESF system in the k-th iteration.

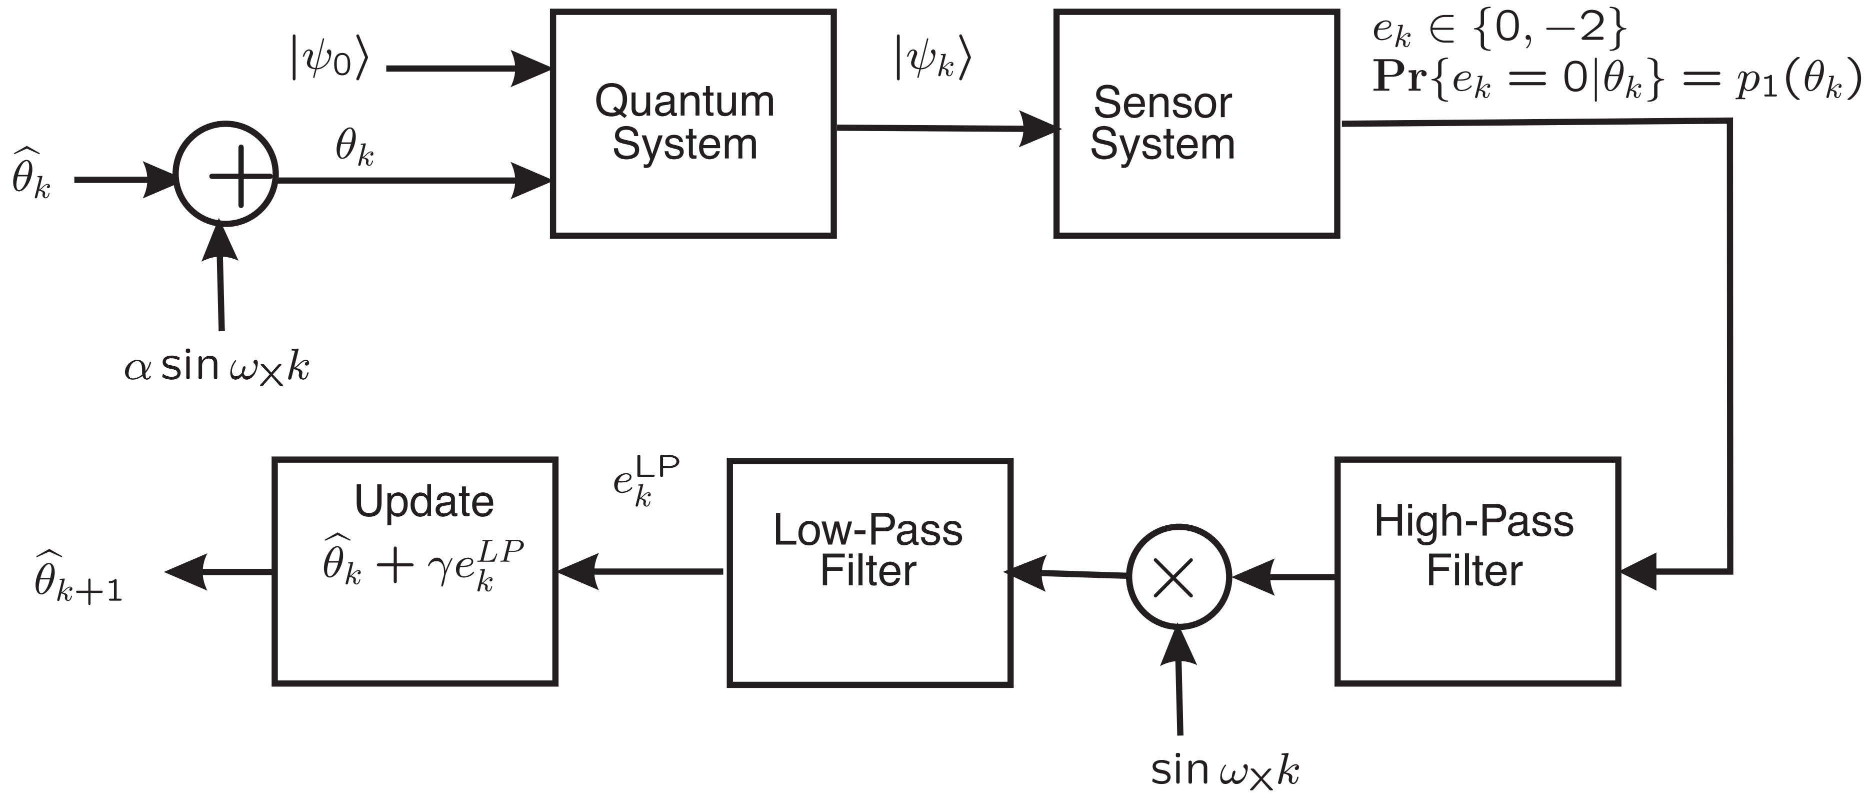

"Theory" IF:

- $< e_k > \approx -2 - \left(\frac{\partial^2 p_1}{\partial \theta_k^2}(\theta_\star)/2\right)(\theta_k - \theta_\star)^2$ with $\frac{\partial^2 p_1}{\partial \theta_k^2}(\theta_\star) < 0$
- $\omega_{\mathrm{LP}} \leq \omega_{\mathrm{HP}} \leq \omega_{\mathrm{X}}$
- $\gamma\alpha\frac{\partial^2 e_k}{\partial \theta_k^2}(\theta_\star) < 0$ and "small"

THEN: $\widehat{\theta}_k \to \theta_\star$

* K.B. Ariyur & M. Krstic, *Real-Time Optimization by Extremum Seeking Feedback*, Wiley, 2003.

With $\gamma = 0.25$, $\alpha = 0.025$, (A), (B), and (C) shows, respectively, typical responses of 300 iterations from the same initial control setting $\theta_0 = 1$.

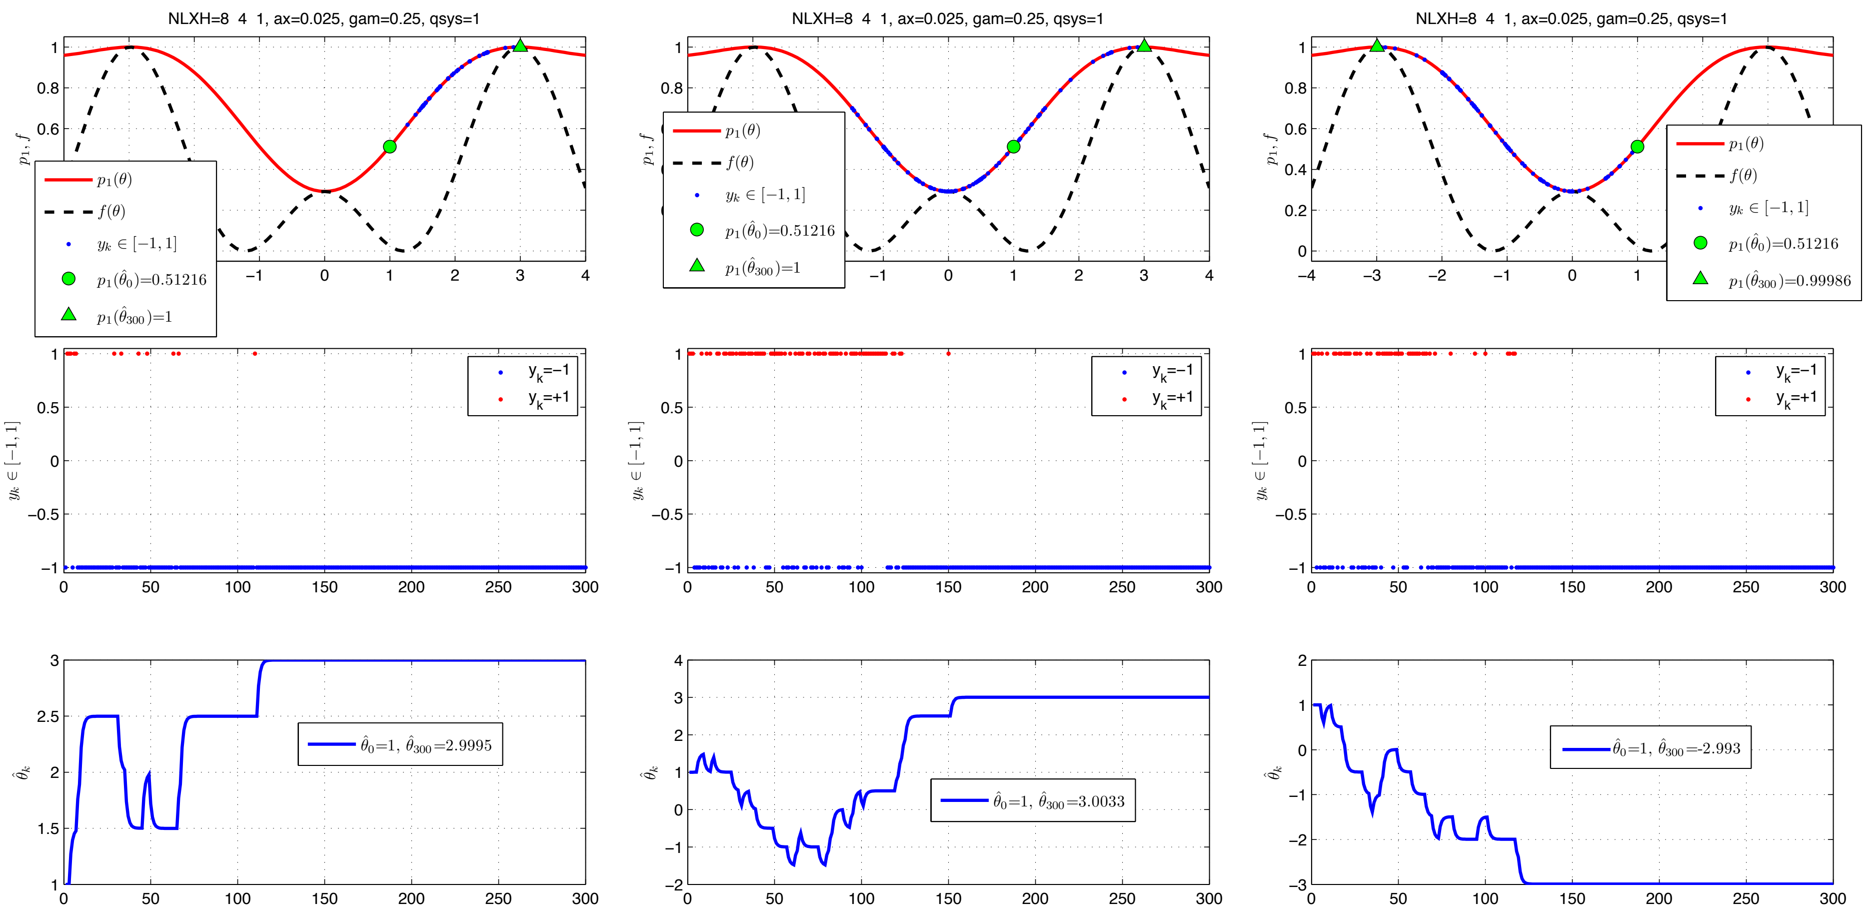

(A) (B) (C)

Extremum Seeking Feedback*

ESF system in the k-th iteration.

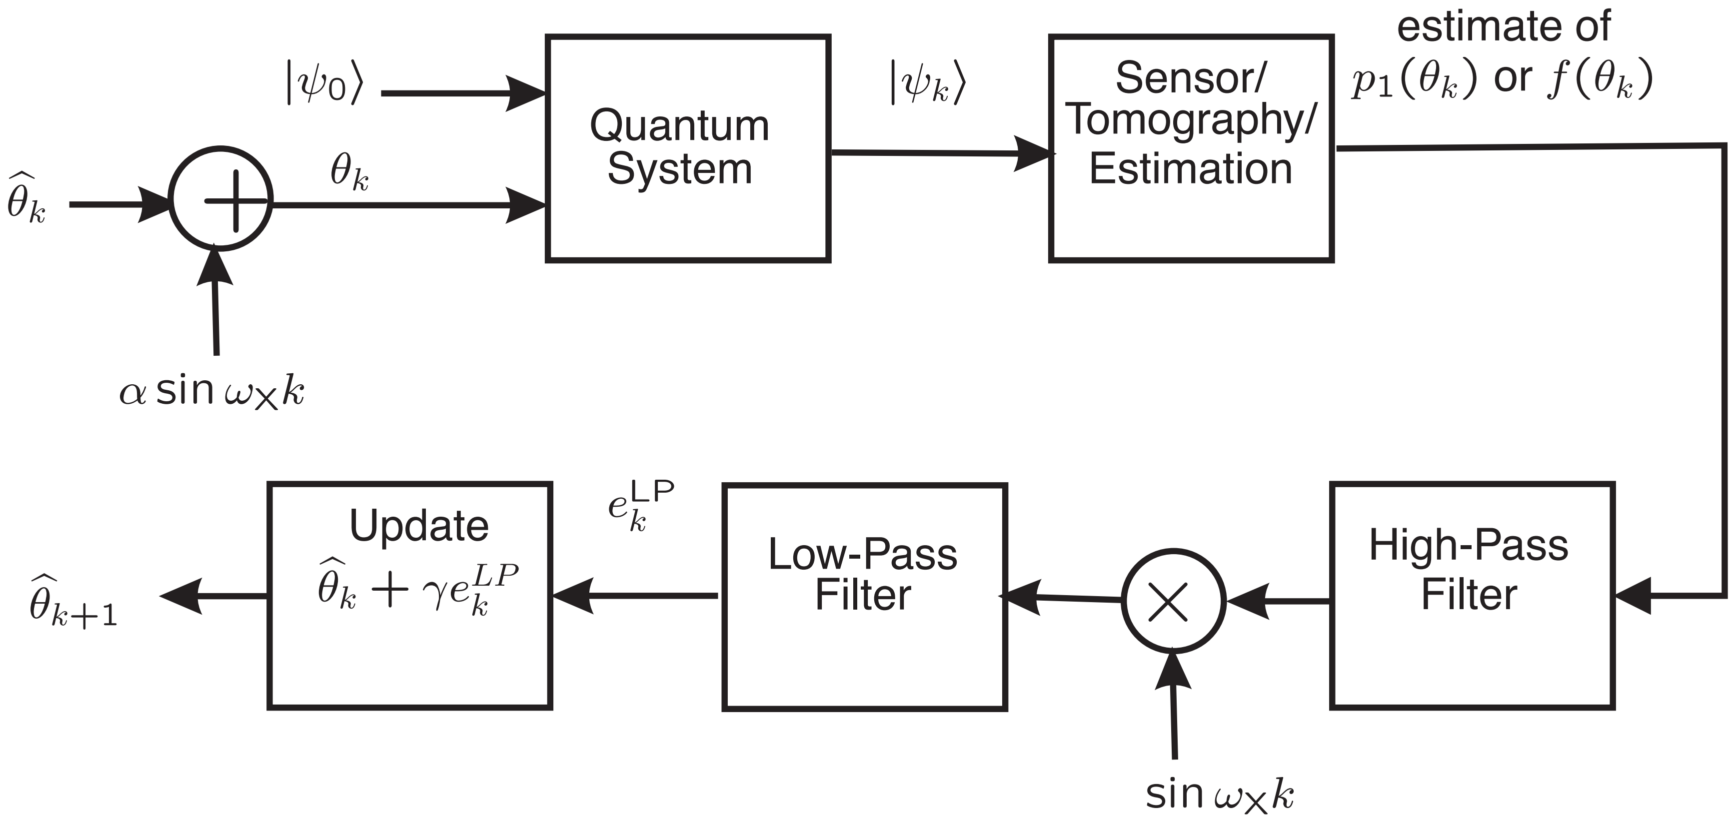

"Theory" IF:

- $f(\theta_k) \approx 1 + \left(\frac{\partial^2 f}{\partial \theta_k^2}(\theta_\star)/2\right)(\theta_k - \theta_\star)^2$ with $\frac{\partial^2 f}{\partial \theta_k^2}(\theta_\star) < 0$
- $\omega_{LP} \leq \omega_{HP} \leq \omega_X$
- $\gamma\alpha\frac{\partial^2 f}{\partial \theta_k^2}(\theta_\star) < 0$ and "small"

THEN: $\widehat{\theta}_k \to \theta_\star$

* K.B. Ariyur & M. Krstic, *Real-Time Optimization by Extremum Seeking Feedback*, Wiley, 2003.

With $\gamma = 0.25$, $\alpha = 0.025$, (A), (B), and (C) shows, respectively, typical responses of 300 iterations.

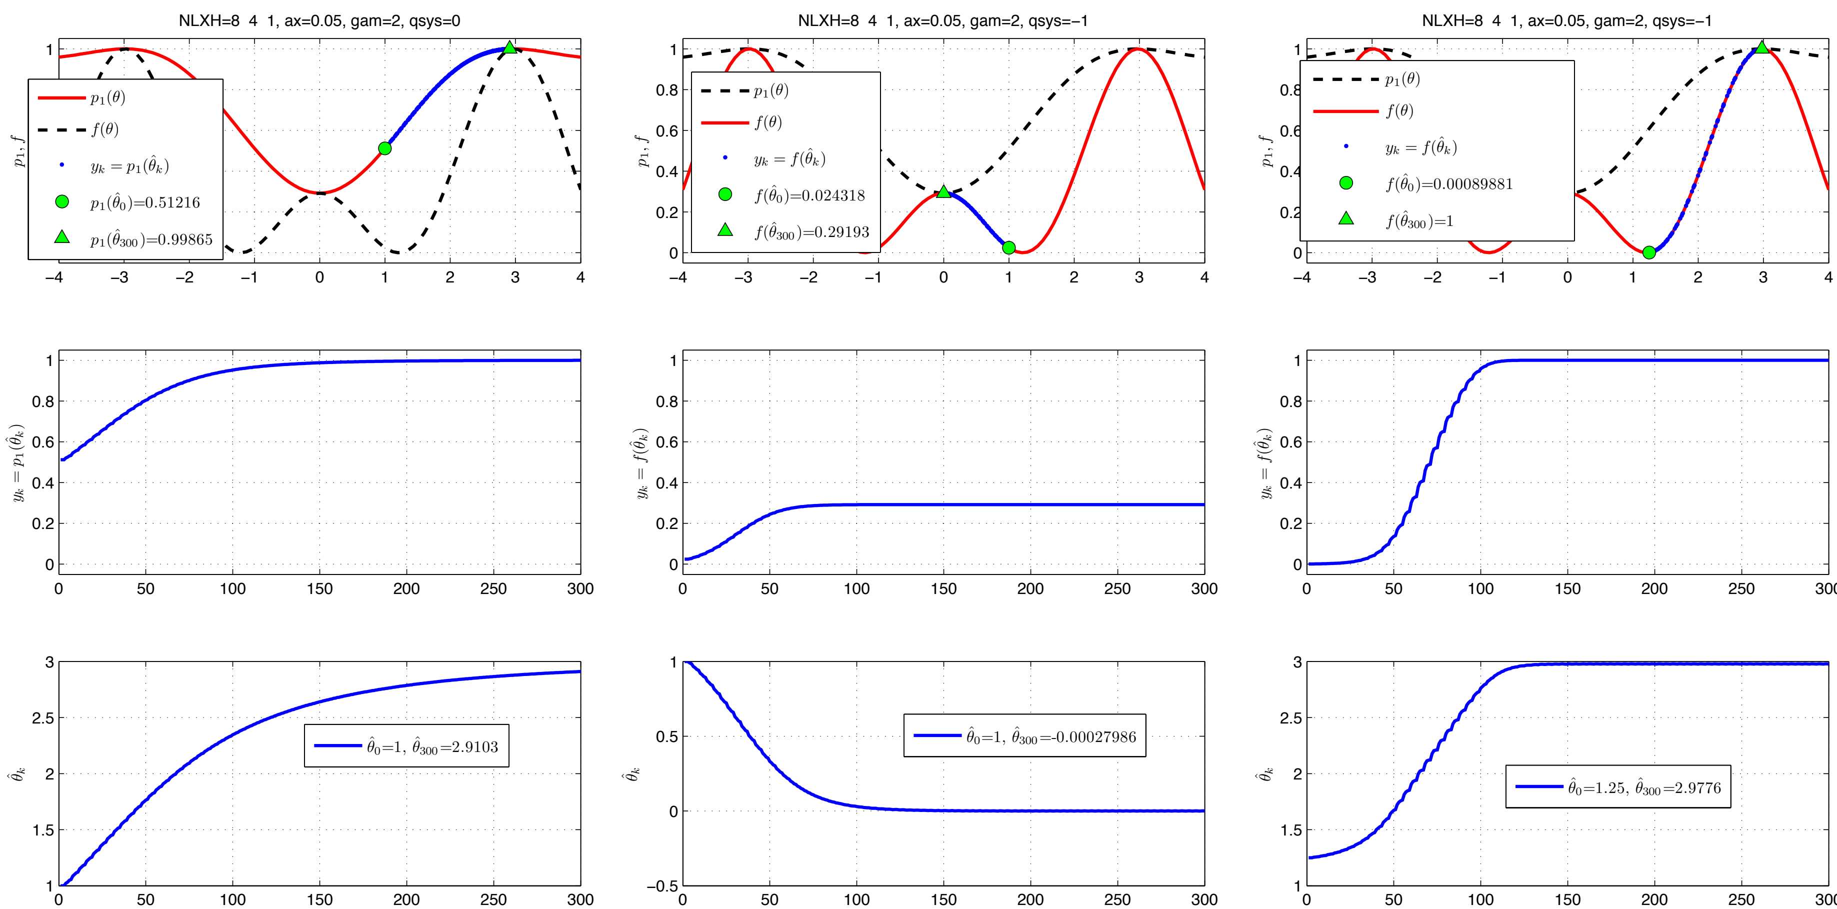

(A)	(B)	(C)
$\theta_0 = 1$	$\theta_0 = 1$	$\theta_0 = 1.25$
estimating $p_1(\theta)$	estimating $f(\theta)$	estimating $f(\theta)$

Summary

- preliminary results are encouraging
  - theory extends to multi-parameter adaptation
  - filters and probing signal can be selected based on level of prior system knowledge, *e.g.*, curvature of map from control to sensor outcome?
- adaptive/learning for discovery, understanding or performance?
  - prior models for each have different levels of detail.
- ESF theory can/may:
  - say something about how to use prior knowledge
  - give lower bounds on data requirements
  - help with sensor selection
- what about ESF for closed-loop (real-time) control?

Extremum Seeking Feedback: Closed-Loop

Quantum system in the k-th iteration:

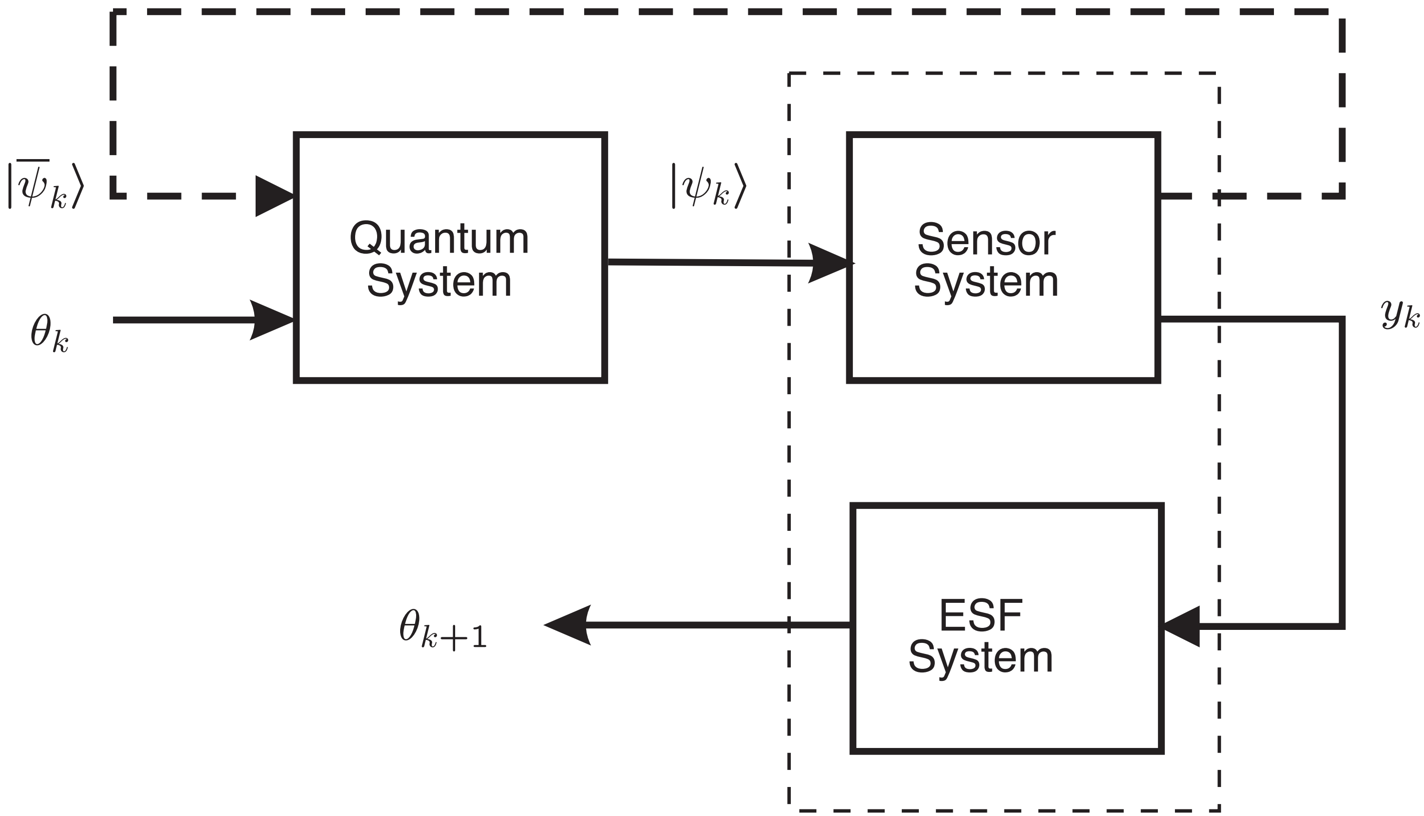

- initial state $|\overline{\psi}_k\rangle$ depends on previous measurement outcome.
- $\theta_k \in \mathbf{R}^n$ are parameters that define k-th control.
- sensor and feedback systems inside the dashed box are "classical".
- sensor system is measurement probe/meter together with data processing capabilities.
- feedback (adaptive/learning) system generates the control θ_{k+1} for the next iteration based on past measurements and controls.

The Feedback Control design paradigm

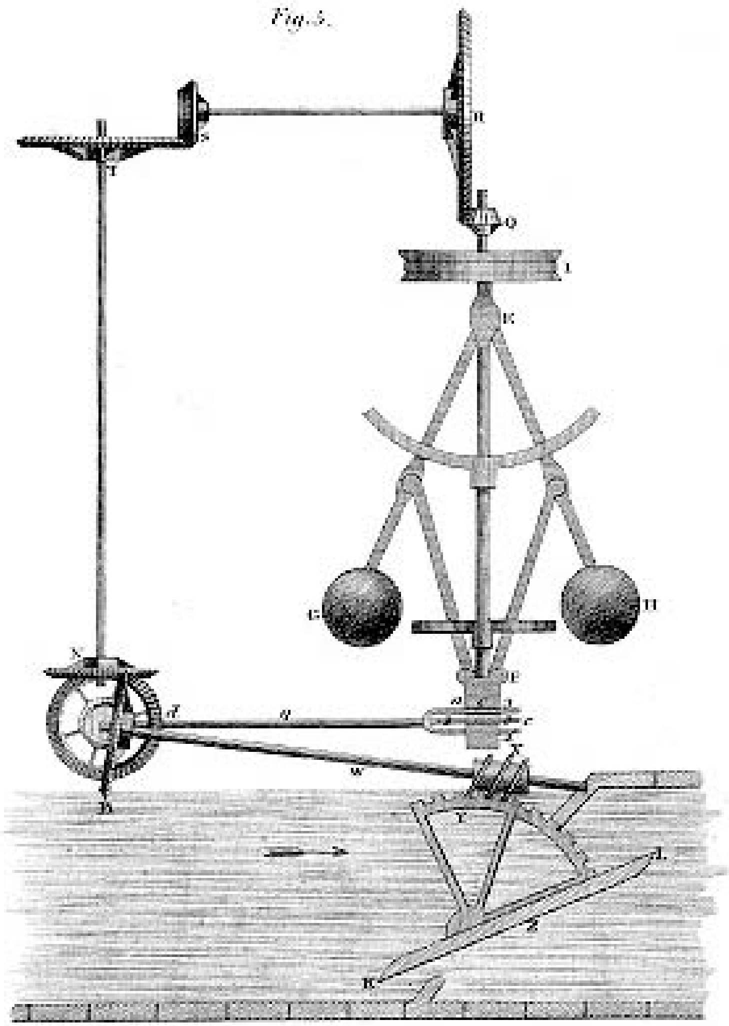

“Flyball Governor”
for steam engine control
James Watt (ca. 1788)

completely mechanical
“natural” speed regulation

Is there a
Quantum Flyball Governor?

completely quantum mechanical
“natural” quantum error correction