Optimal Directed Search on Rugged Landscapes

August 2013

Steven Callander
Stanford University

Niko Matouschek
Northwestern University

Motivation

- Large and persistent performance differences.
- Implications for macro, IO, labor, etc.
- But why do performance differences arise in the first place?
- We explore differences in managerial practices.

Motivation

- Do managerial practices affect performance?
- But why do managers adopt different practices?
- And why don't underperformers imitate top performers?

Motivation

- Our premise:
  - Managerial problem too complex to be solved analytically.
  - Instead, managers learn by trial-and-error.
- Our main results:
  - Learning by trial-and-error gives rise to performance differences.
  - Learning by trial-and-error + complementarities + a little bit of firm-specificity can make imitation very risky.

Agenda

- The Model
- Beliefs and Expected Utility
- Managerial Learning
- Persistent Performance Differences
- Barriers to Imitation
- Conclusions

The Model

- One manager.
- One managerial practice that can take any value on the real line.
- Time $t = 1, 2, \ldots$
- In any period *t*, the manager
  - takes action a_t
  - which generates income $m_t = m(a_t)$
  - which, in turn, gives the manager utility $u(m_t)$.

The Production Function

- The production function is the realized path of a Brownian motion
  - with drift ❍ > 0 and variance ♦$^2 > 0$
  - in which the action replaces time as the independent variable.

The Production Function

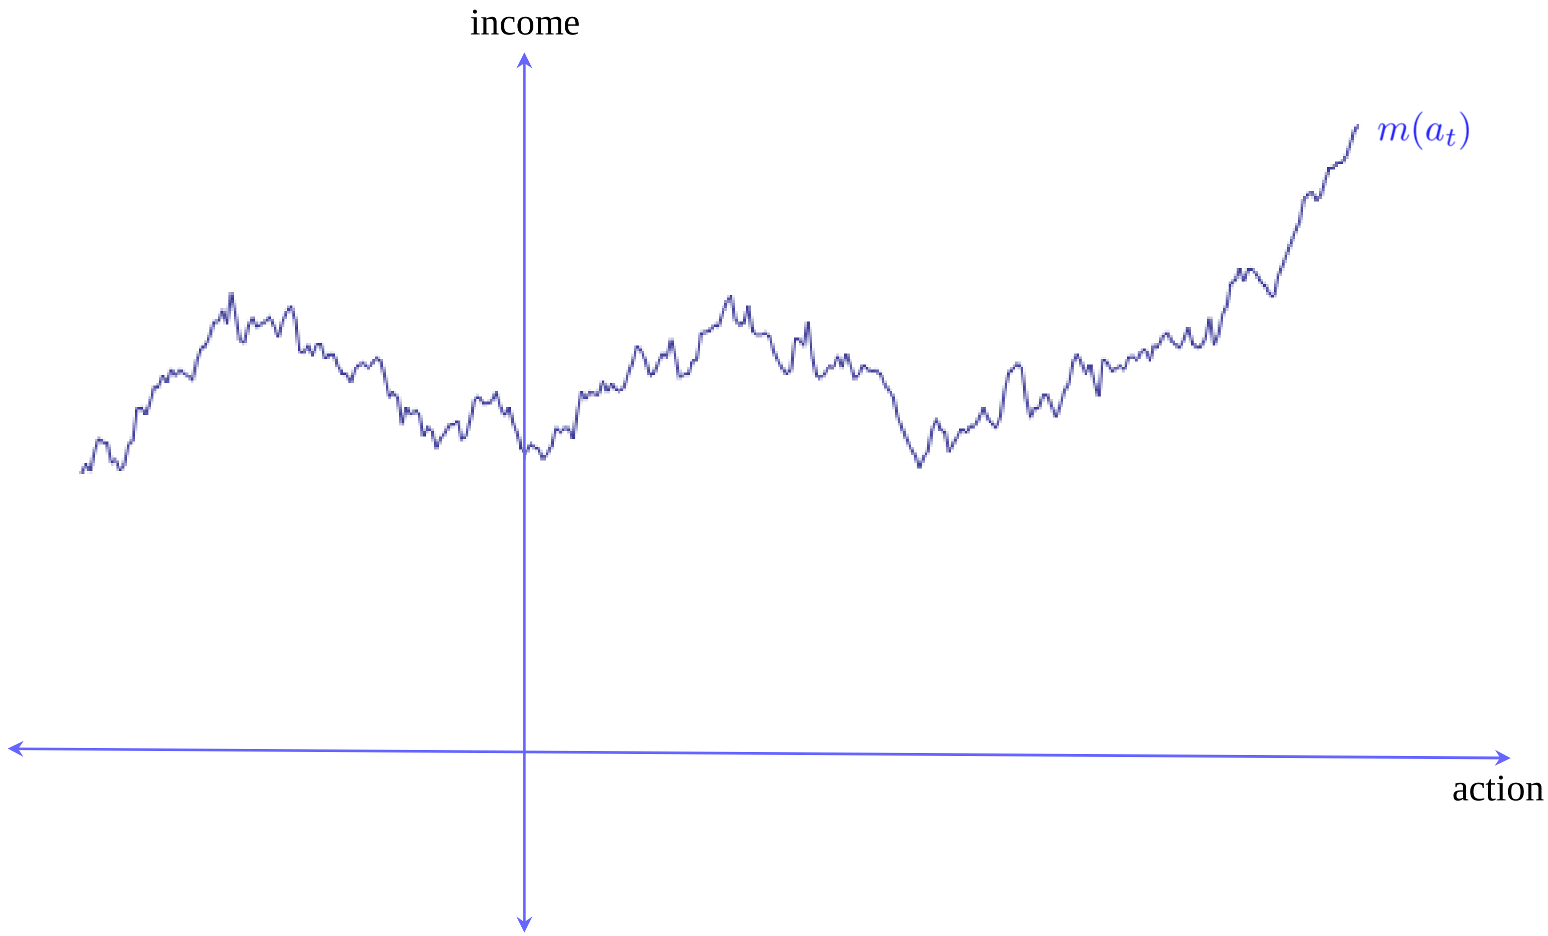

The Production Function

- The production function is the realized path of a Brownian motion
  - with drift ❍ > 0 and variance ♦$^2 > 0$
  - in which the action replaces time as the independent variable.
- Status quo action a_0 and status quo income $m_0 = m(a_0)$.

The Production Function

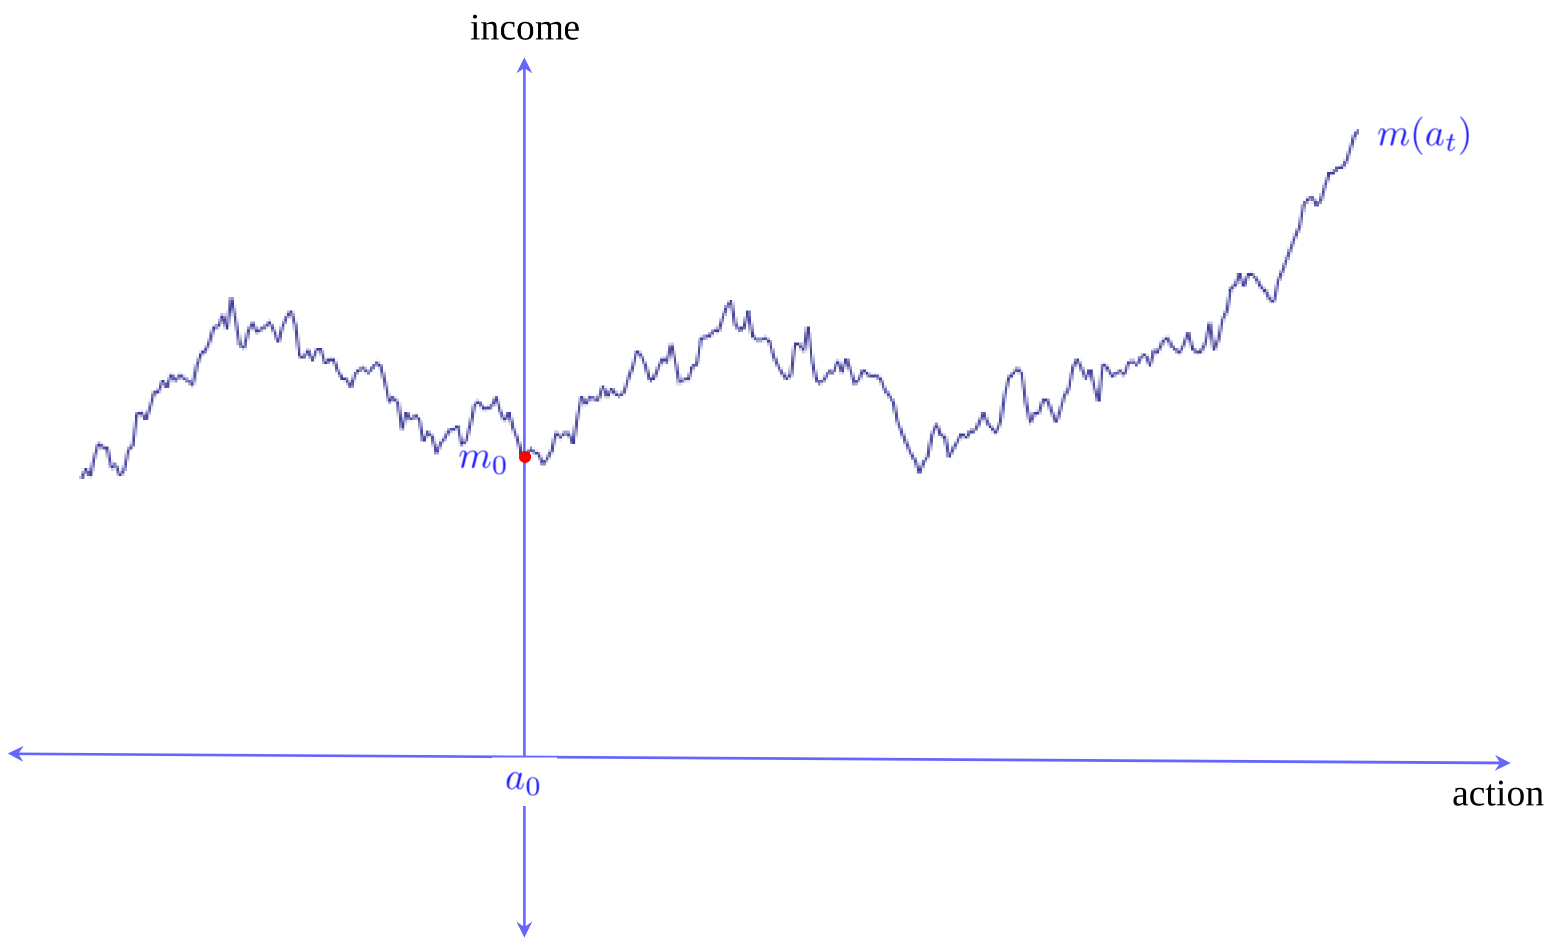

THE PRODUCTION FUNCTION

- The production function is the realized path of a Brownian motion
  - with drift ❍ > 0 and variance ♦$^2 > 0$
  - in which the action replaces time as the independent variable.
- Status quo action a_0 and status quo income $m_0 = m(a_0)$.
- The variance ♦^{2} measures the complexity of the production process:

  larger ♦^{2} ←→ production more complex ←→ landscape more rugged
- Note that the production function does not change over time.

The Manager's Preferences

- Denote:
  - the utility function by $u(m_t)$,
  - the coefficient of absolute risk aversion by

$$r\left(m_t\right) = -\frac{u''\left(m_t\right)}{u'\left(m_t\right)}$$

  - and the coefficient of absolute prudence by

$$p\left(m_t\right) = -\frac{u'''\left(m_t\right)}{u''\left(m_t\right)}$$

The Manager's Preferences

- Preferences satisfy:
  - Non-satiation: $u'(m_t) > 0$
  - Risk aversion: $u''(m_t) < 0$
  - Decreasing absolute prudence: $p'(m_t) \leq 0$
  - Decreasing absolute prudence implies DARA, i.e., $r'(m_t) \leq 0$.
  - Coefficient of absolute risk aversion crosses $2\mu/\sigma^2$

The Manager's Preferences—An Example

- Exponential utility:

$$u(m_t) = \text{♋} m_t - \exp(\text{-}\text{♌} m_t),$$

where ♋ > 0 and ♌ > 2❍$/$♦^{2} .

- For this utility function
  - prudence is given by $p(m_t) = \beta$
  - and risk aversion is given by $r(m_t) = \dfrac{\beta^2 \exp(-\beta m_t)}{\alpha + \beta \exp(-\beta m_t)}$

Information

- The manager does not know the realized path of the Brownian motion.
- But he does know
  - the properties of the Brownian motion: drift ❍ and variance ♦^{2},
  - the status quo action a_0 and associated income m_0, and
  - his previous actions and associated income levels.

Equilibrium

- The manager maximizes expected utility on a period-by-period basis.
- An equilibrium is given by $(a_1^*, a_2^*, ...)$, where

$$a_t^* \in \arg\max_{a_t} \mathrm{E}\left[u\left(m_t\right)|\text{period } t \text{ information}\right]$$

Agenda

- The Model
- Beliefs and Expected Utility
- Managerial Learning
- Persistent Performance Differences
- Barriers to Imitation
- Conclusions

Beliefs

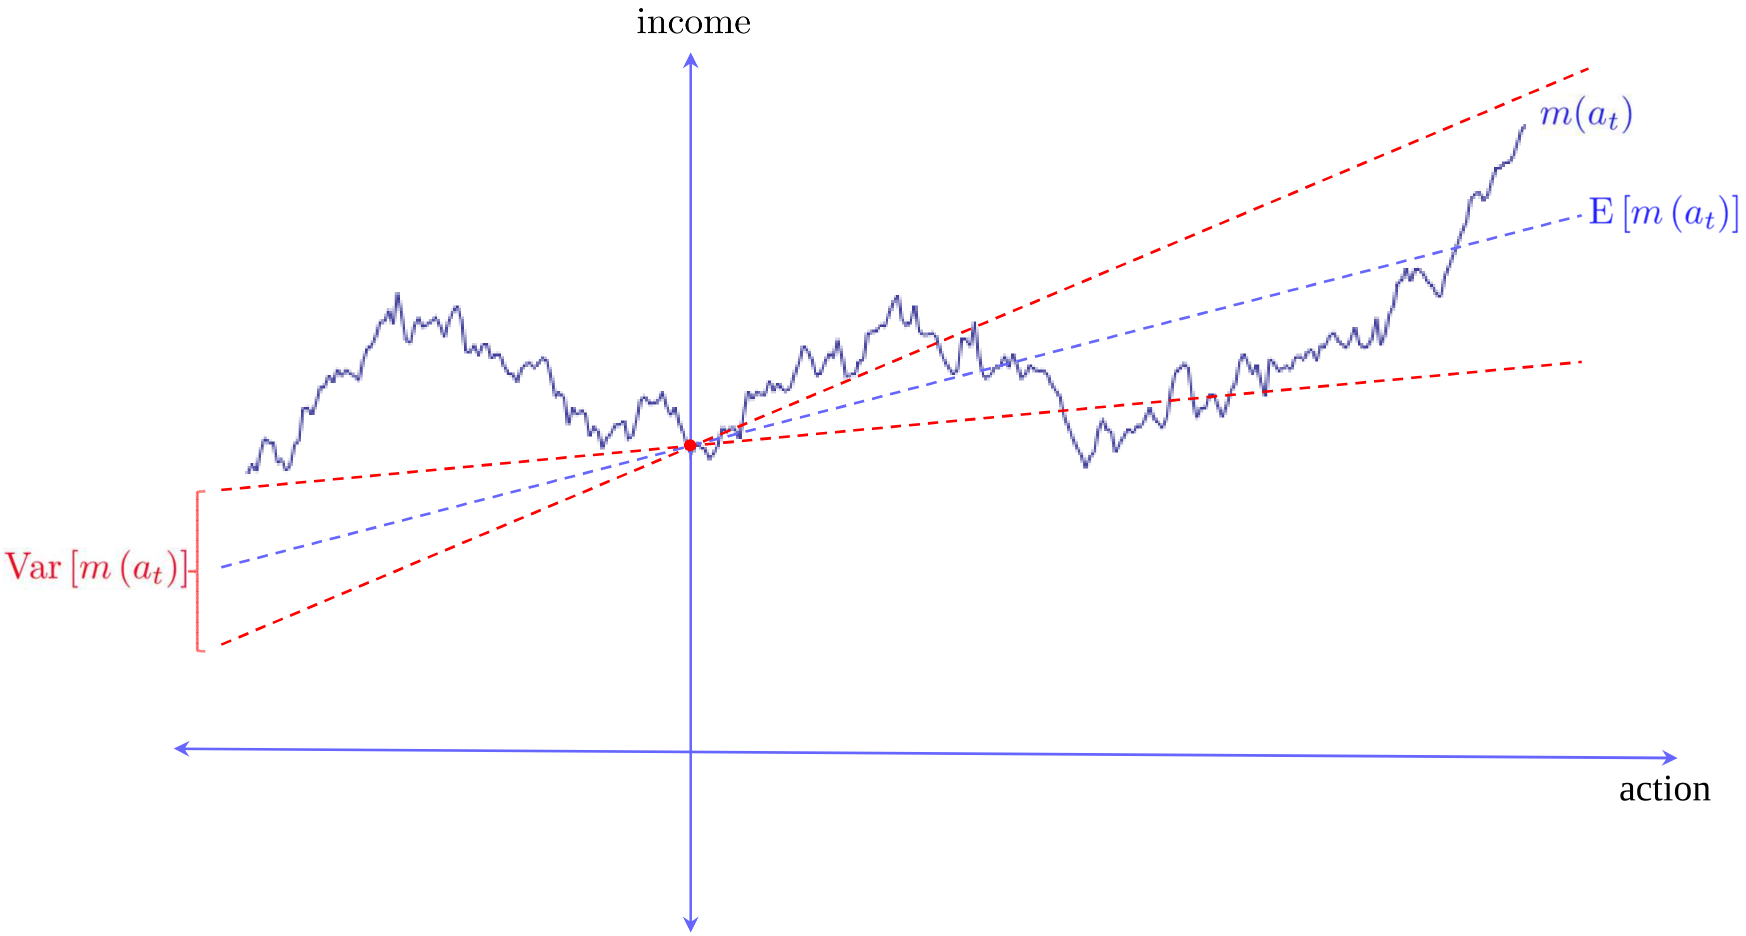

income
$m(a_t)$
$\mathrm{E}\left[m\left(a_t\right)\right]$
$\mathrm{Var}\left[m\left(a_t\right)\right]$
action

Beliefs

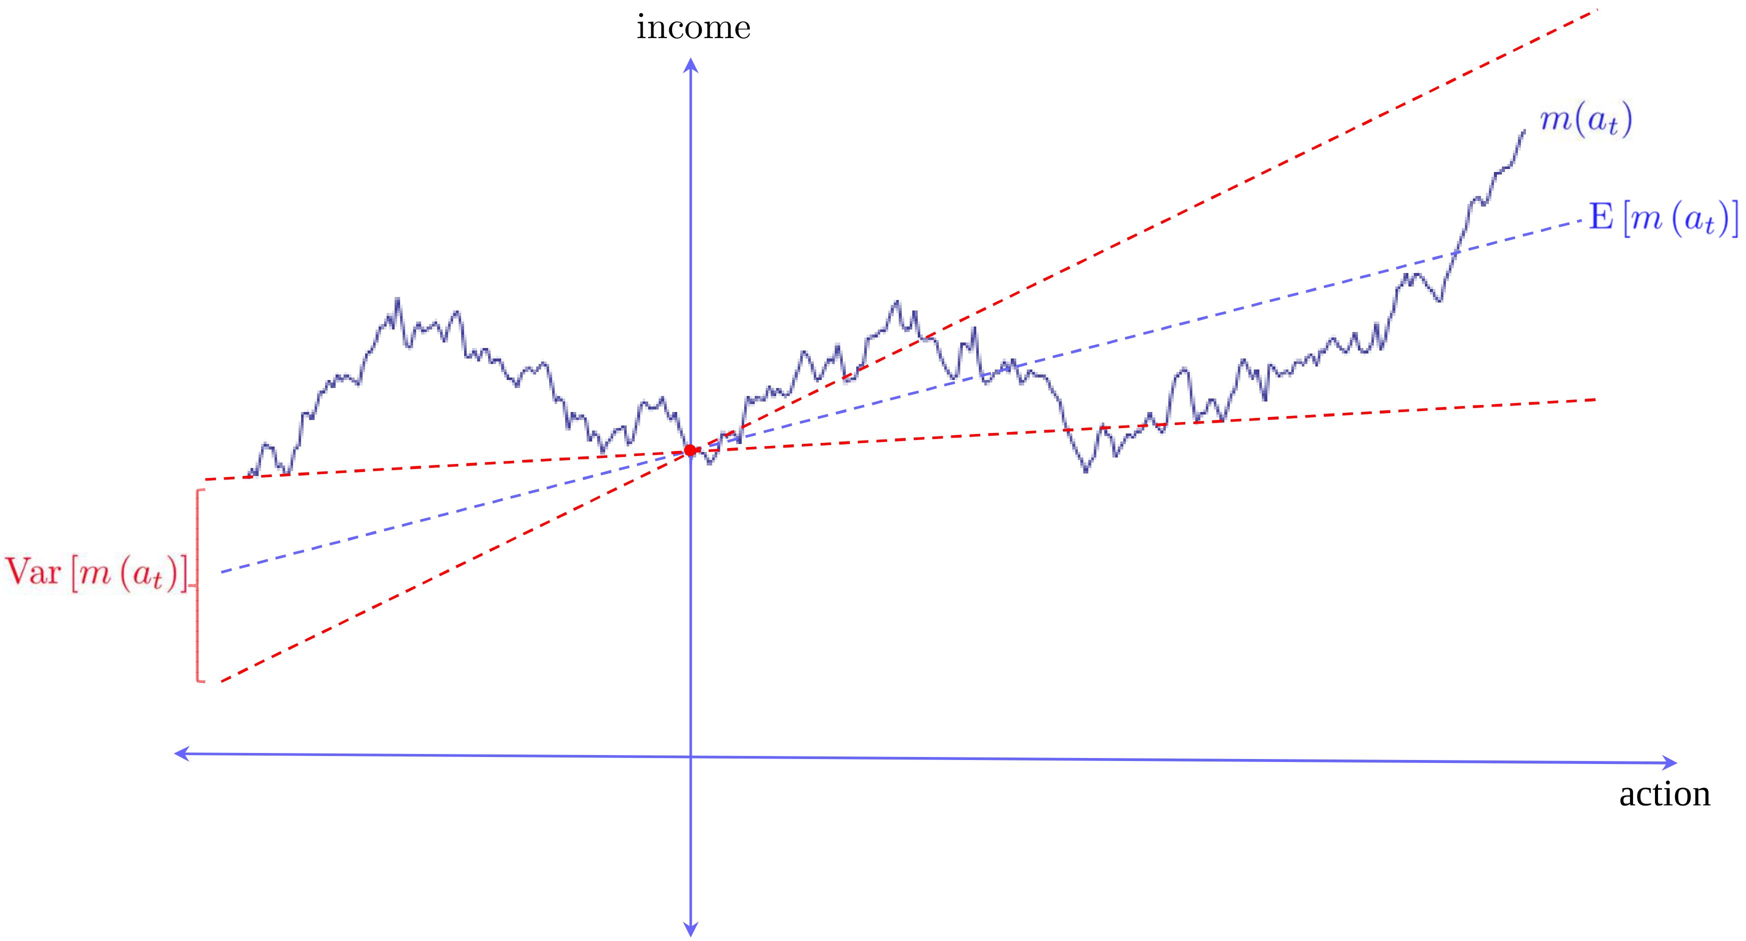

Beliefs

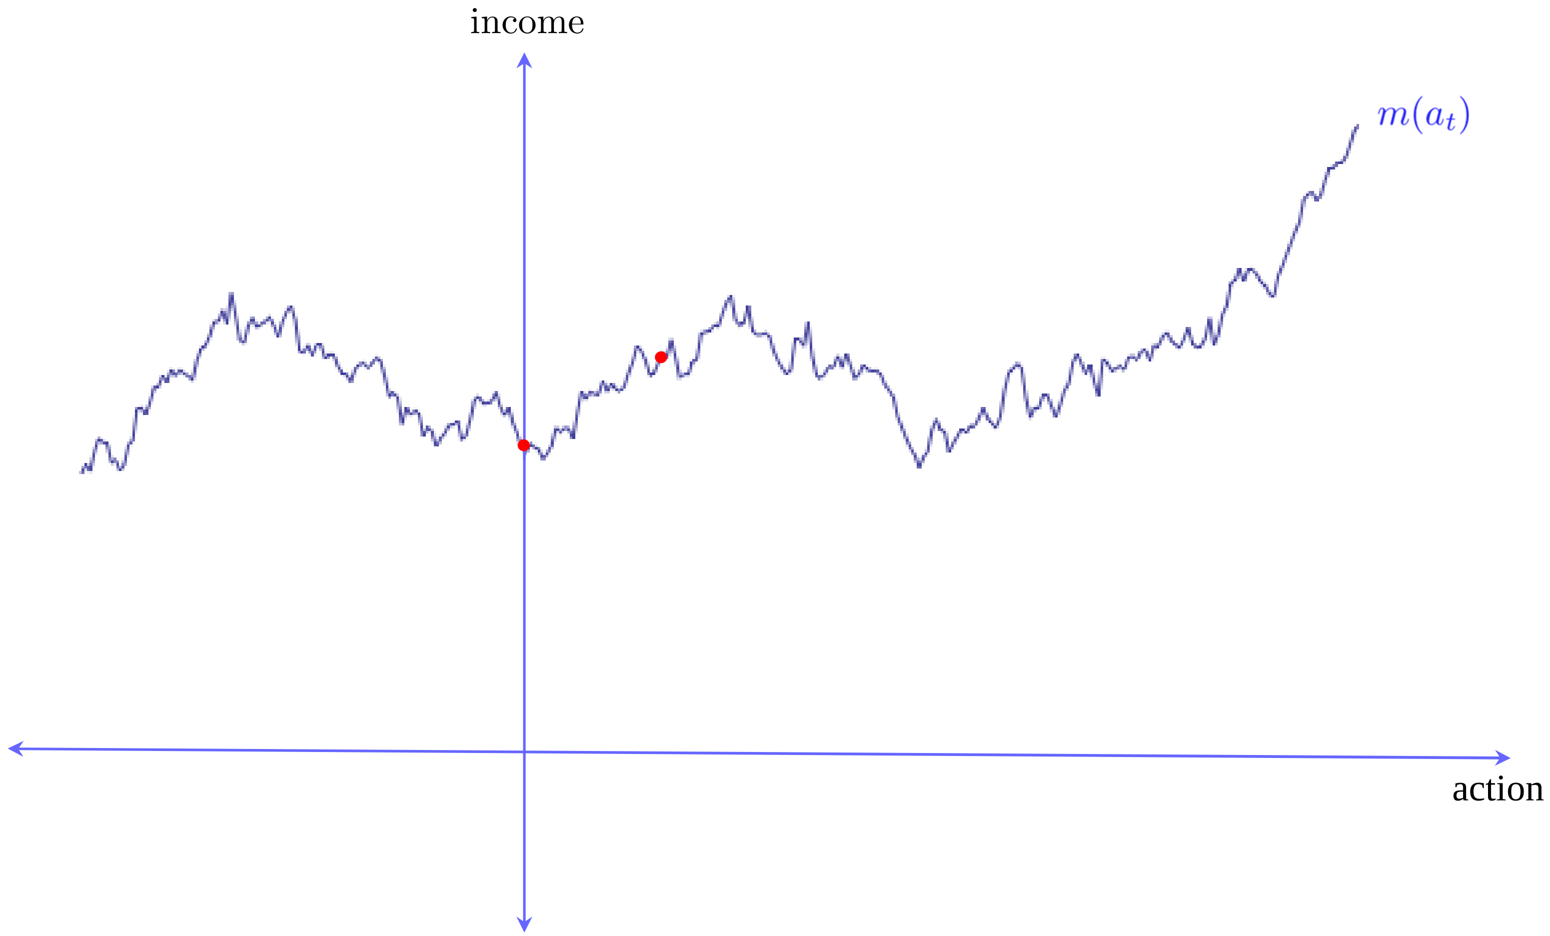

Beliefs

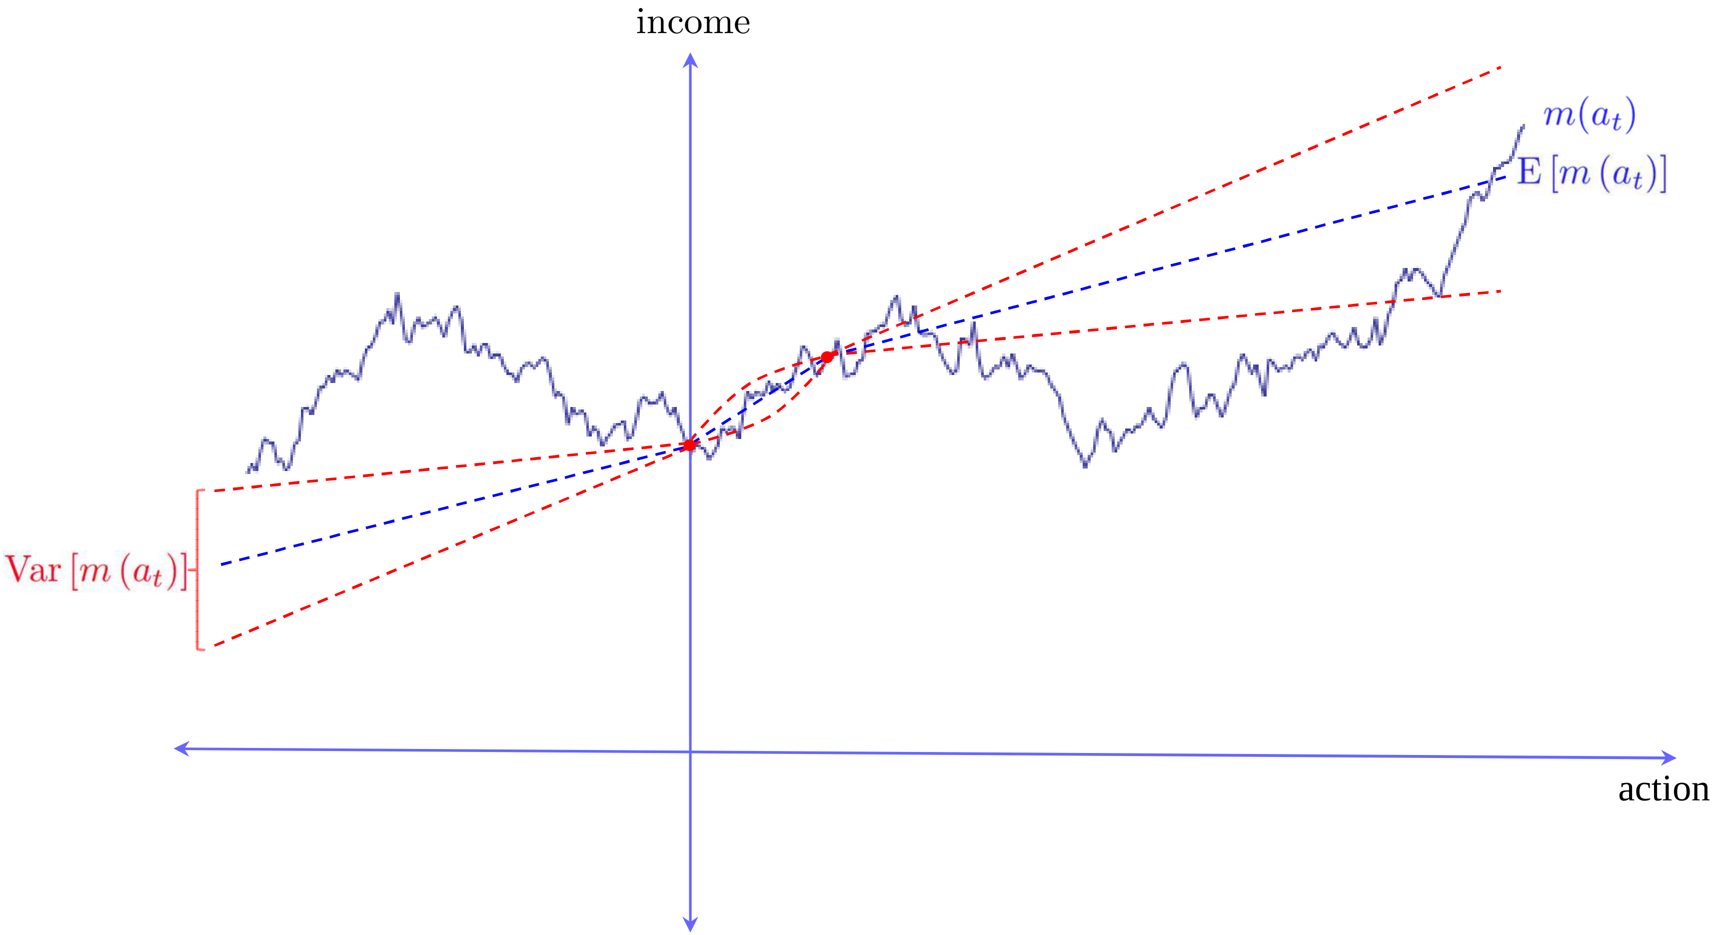

Beliefs

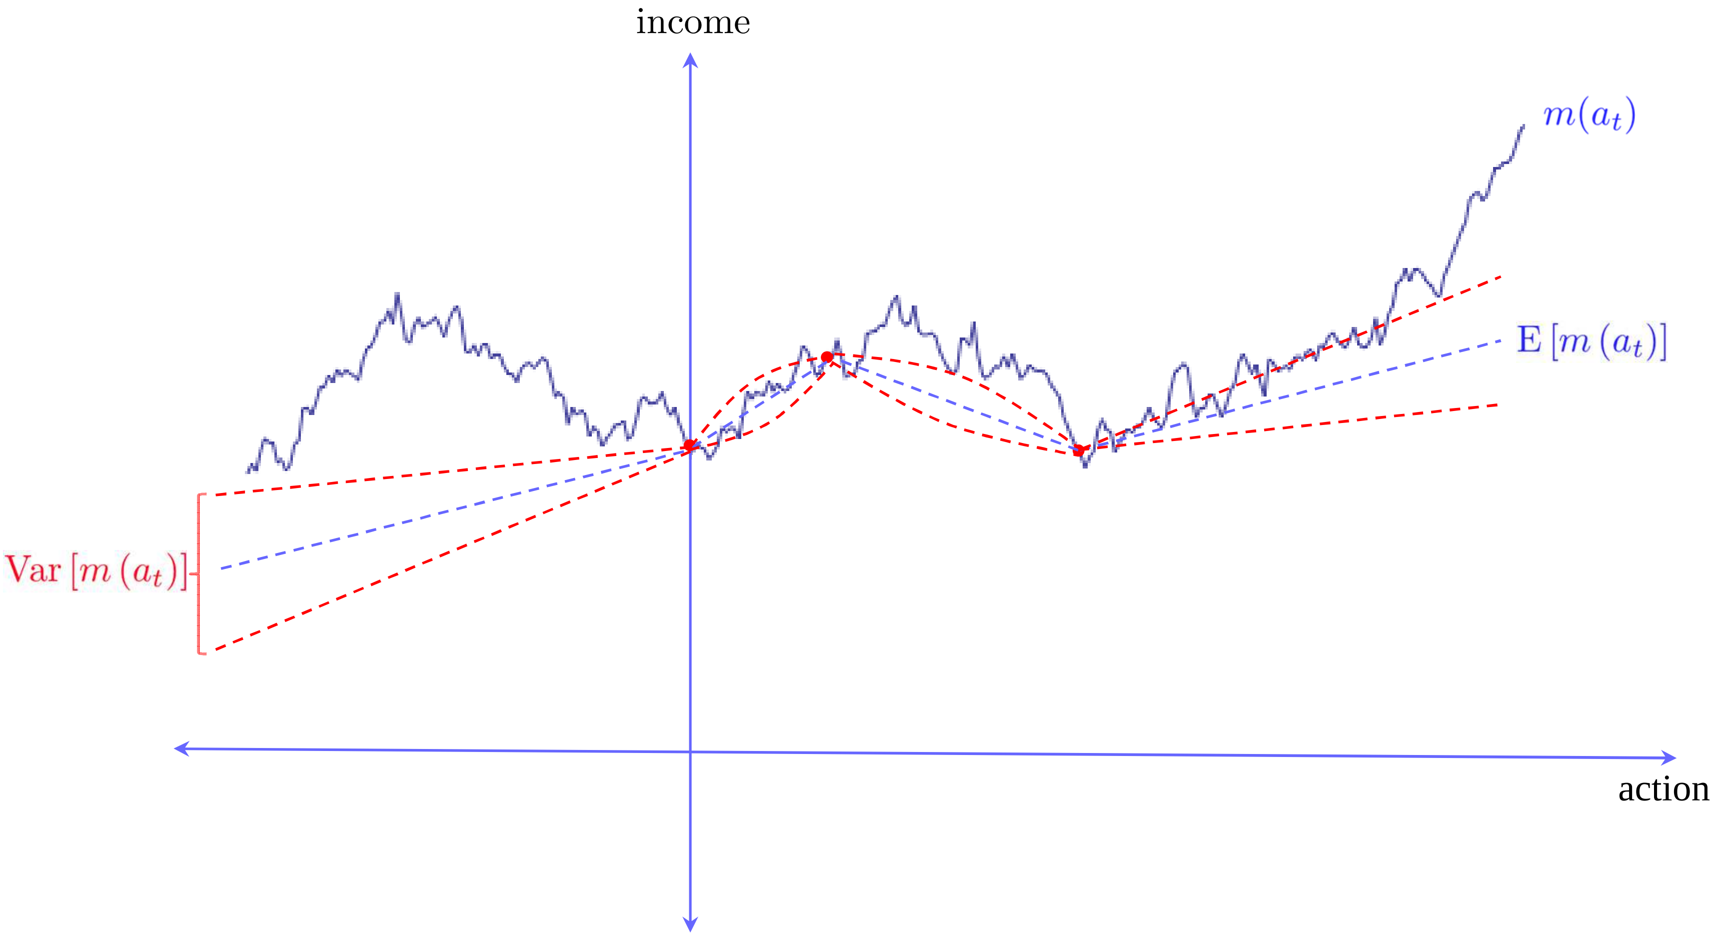

income
$m(a_t)$
$\mathrm{E}\left[m\left(a_t\right)\right]$
$\mathrm{Var}\left[m\left(a_t\right)\right]$
action

Why the Brownian Motion?

- Beliefs are normally distributed.
- Manager knows in which direction to search for better actions.

  → directed search
- Manager knows more about an action...

  ...the closer it is to a known action...

  ...and the less complex the production process is.
- Manager can never infer entire production function.

  → there is a limit to theoretical knowledge...

  ...and thus a need to learn by trial-and-error.

Expected Utility

- Suppose that $m \sim \mathrm{N}(M, V)$.

Expected Utility

- Suppose that $m \sim \mathrm{N}(M, V)$.
- Chipman (1973) and Lajeri and Nielsen(2000) show that the expected utility function

$$W(M, V) = \mathrm{E}\left[u\left(M + \sqrt{V}z\right)\right]$$

is concave.

because of
decreasing absolute prudence

random variable
drawn from
standard normal

Agenda

- The Model
- Beliefs and Expected Utility
- Managerial Learning
  - The First Period
  - The Second and Subsequent Periods
  - The Optimal Learning Rule
- Persistent Performance Differences
- Barriers to Imitation
- Conclusions

First Period—Beliefs

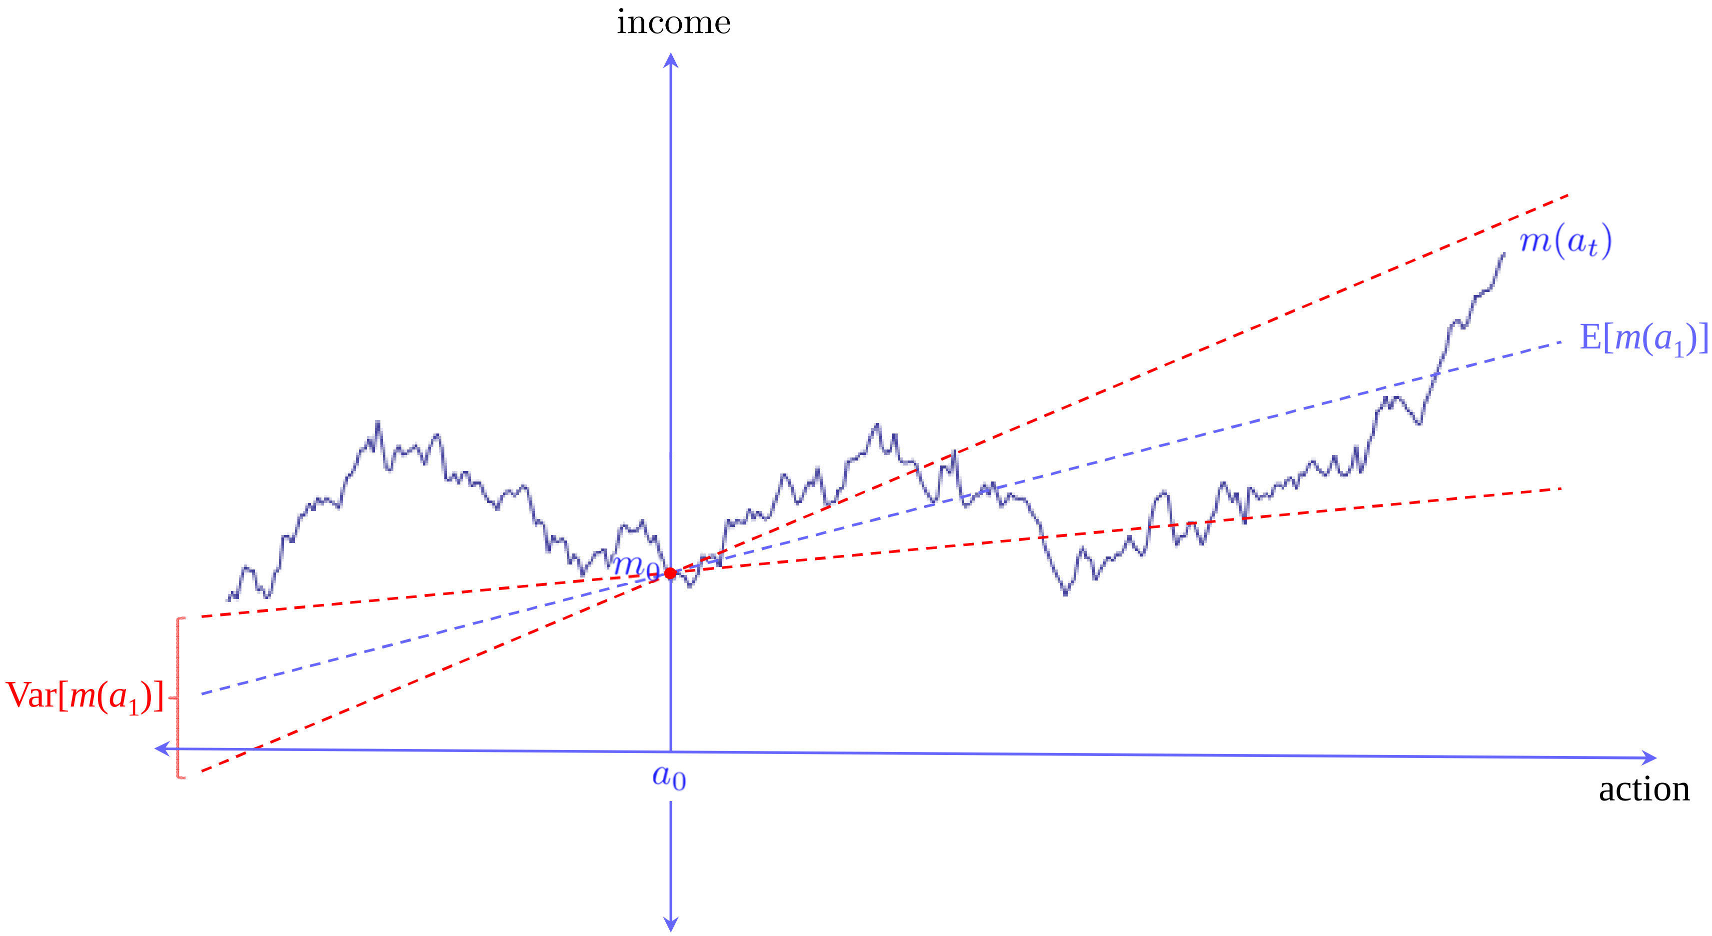

First Period—The Problem

$$\max_{\Delta_1 \geq 0} W\left(m_0 + \mu\Delta_1, \sigma^2\Delta_1\right)$$

$$\Delta_1 = a_1 - a_0 \geq 0$$

First Period—The Problem

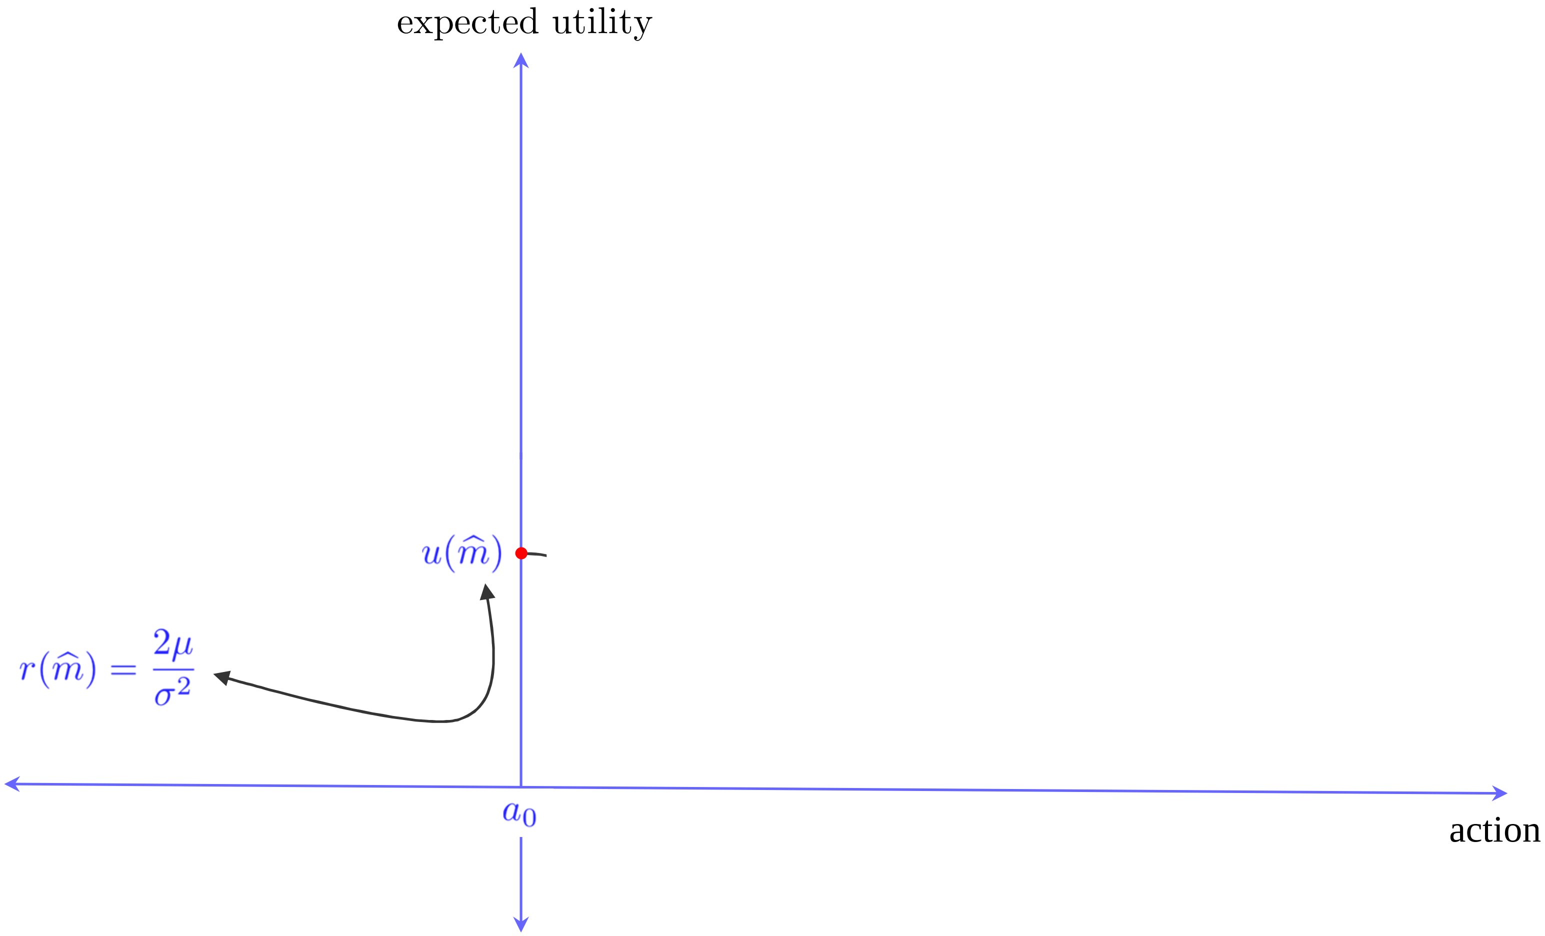

First Period—The Problem

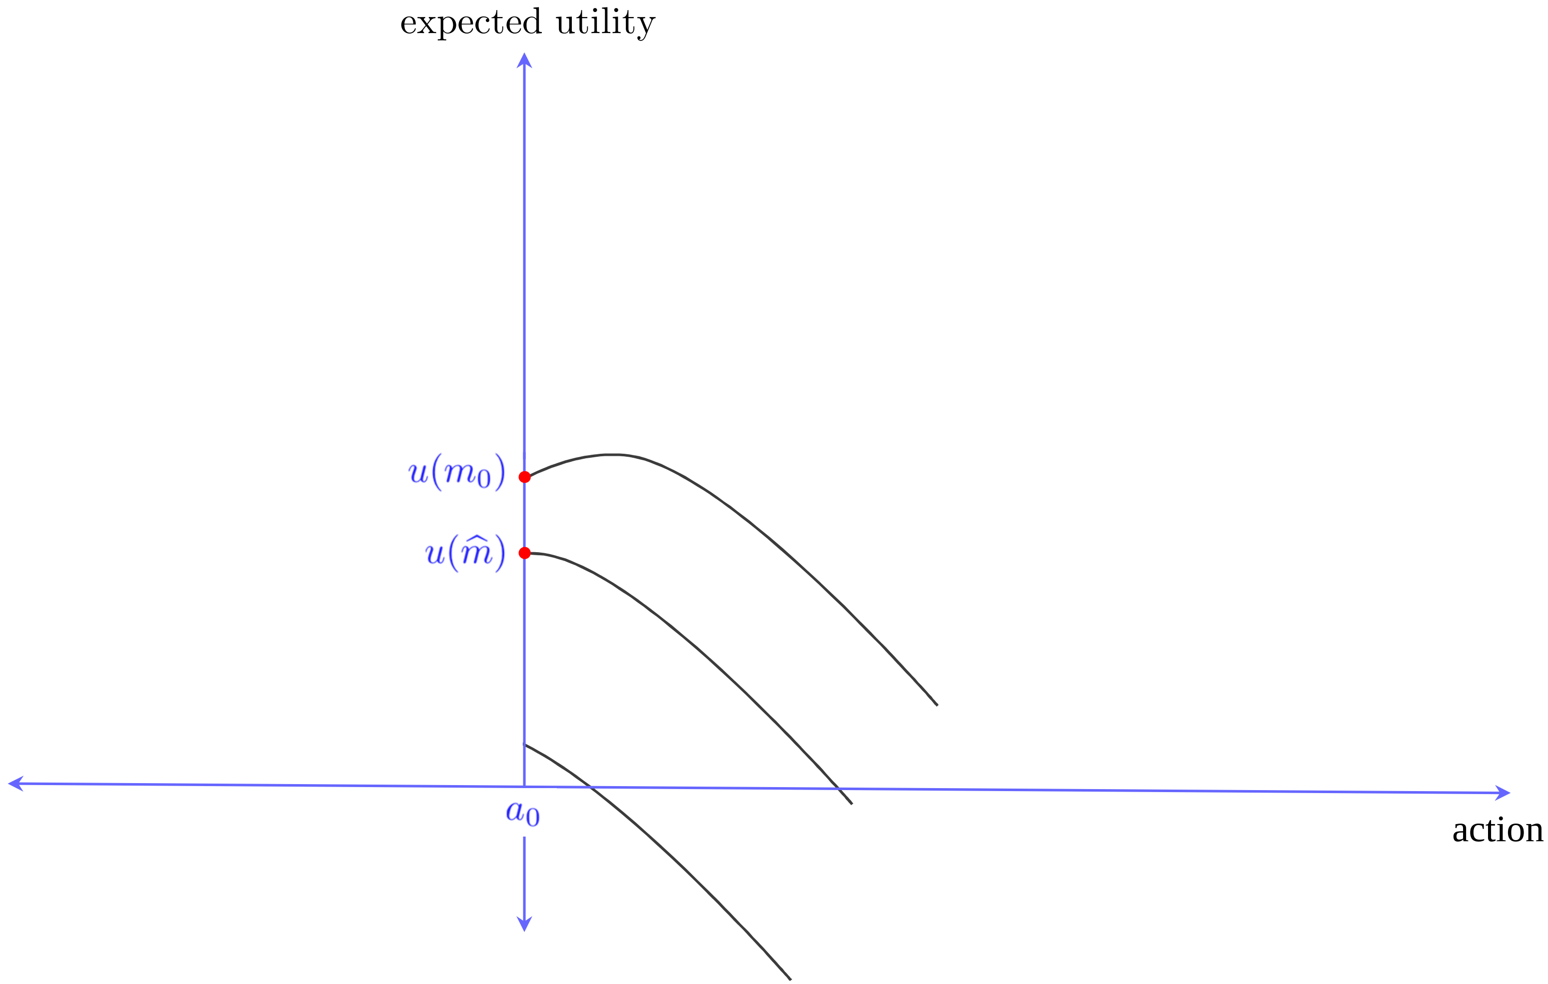

First Period—Optimal Action

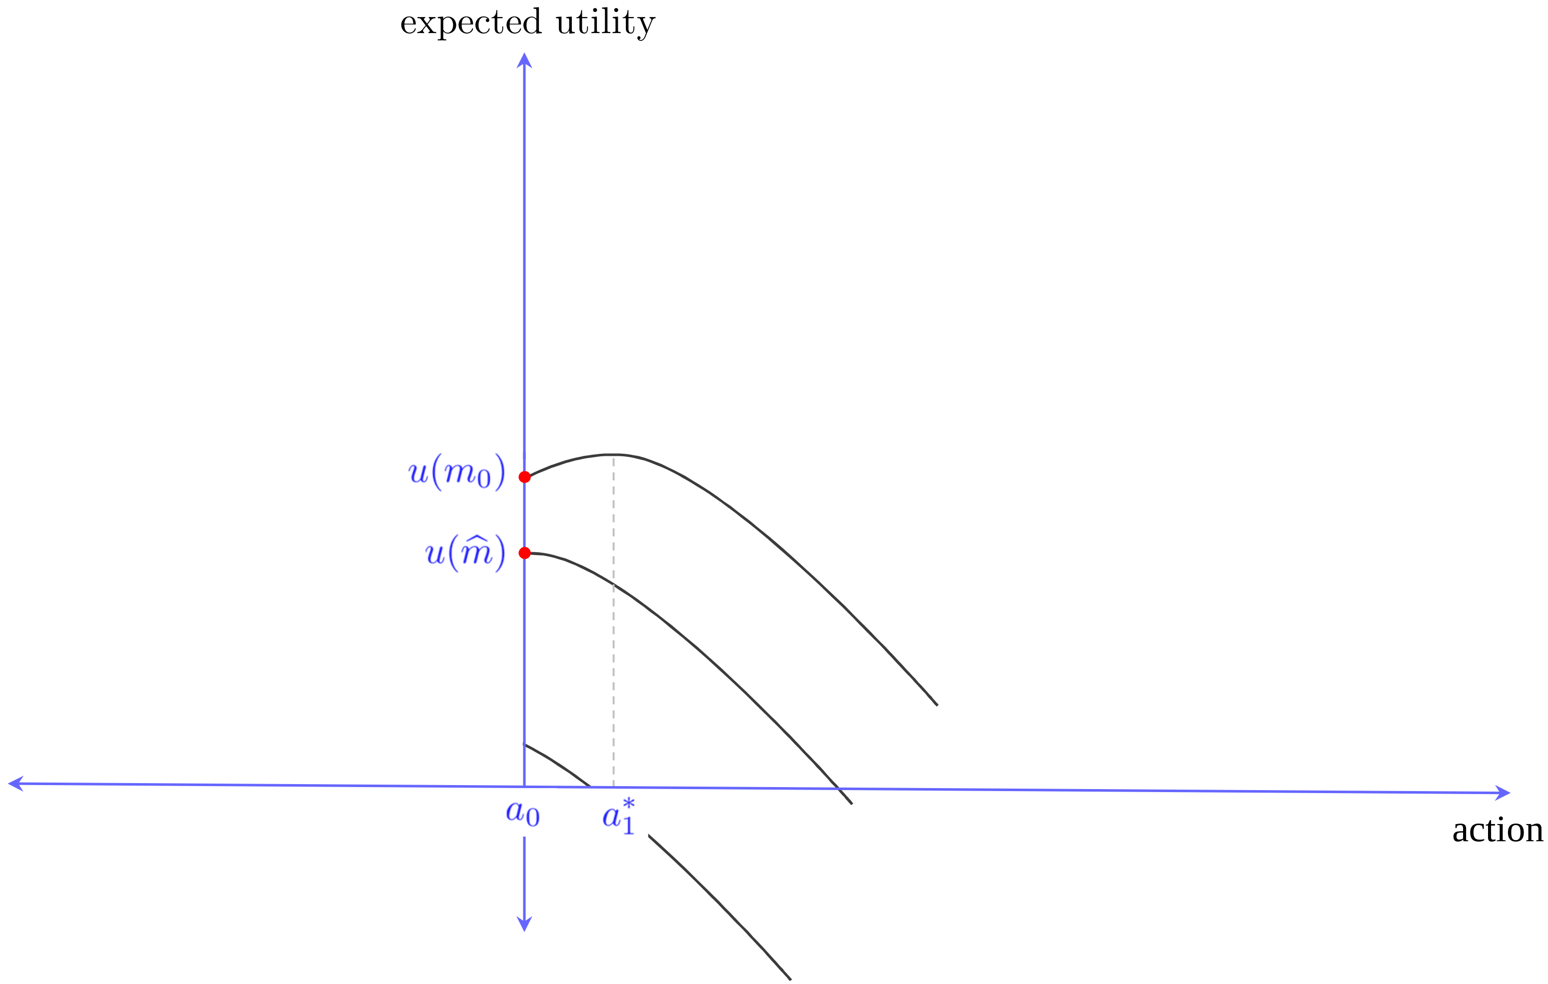

First Period—Optimal Action

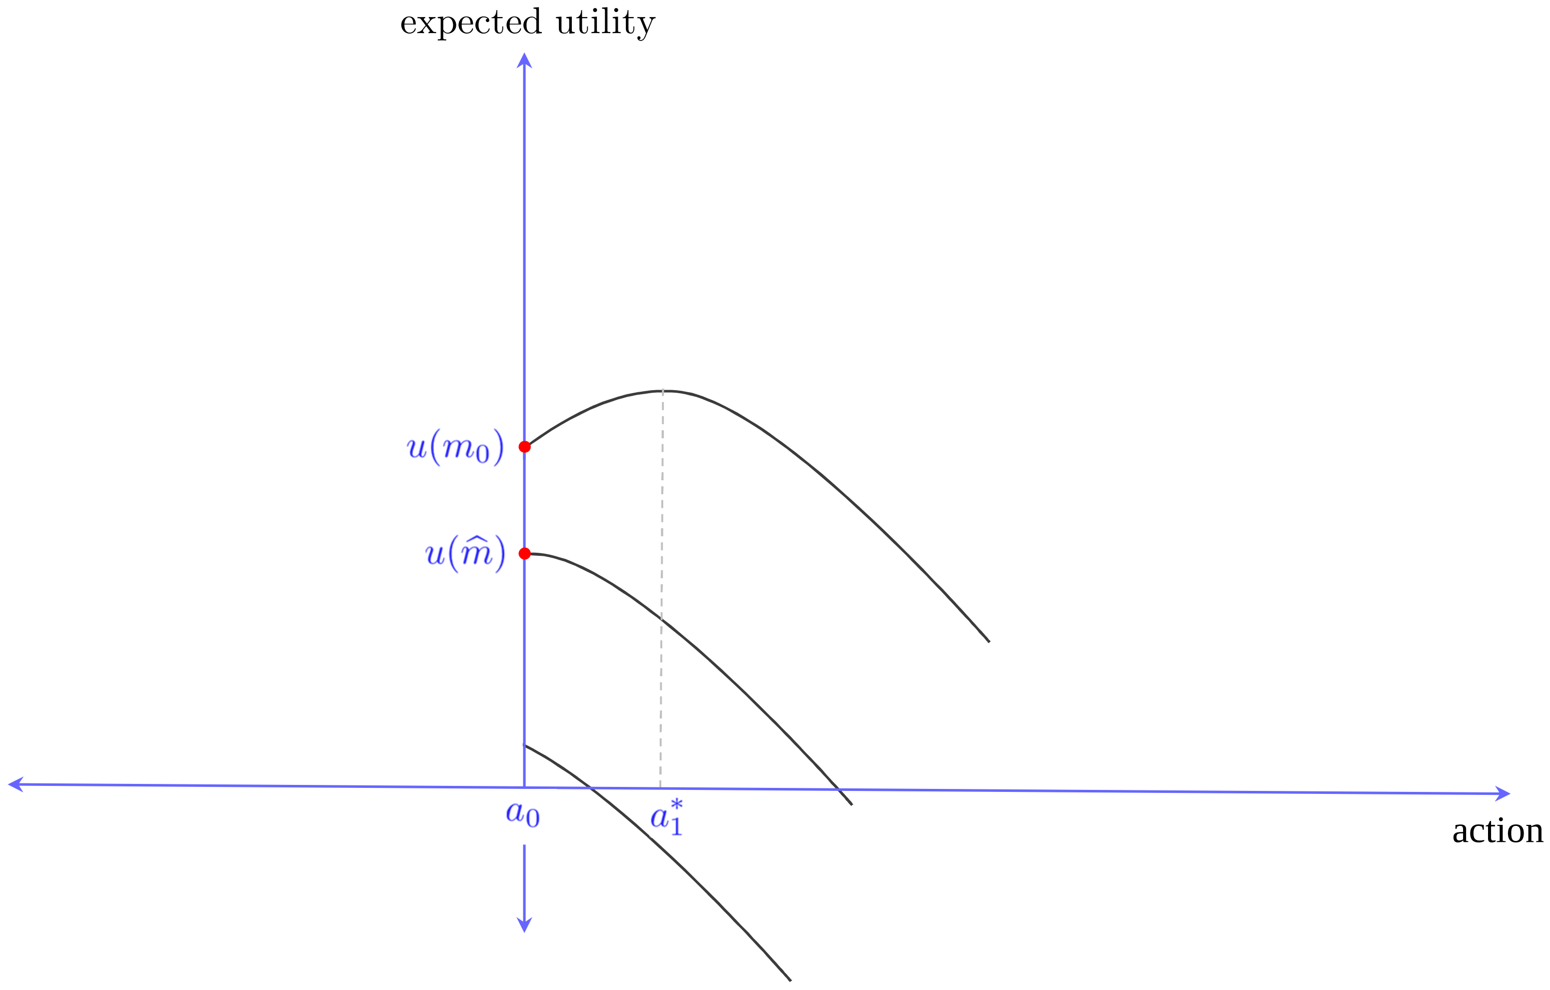

Optimal Learning—First Period

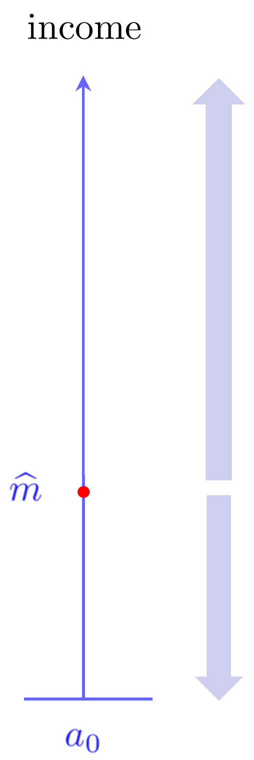

$$r(\widehat{m}) = \frac{2\mu}{\sigma^2}$$
$$\equiv$$
$$-\frac{u''(\cdot)}{u'(\cdot)}$$

take action $a_0 + \Delta(m_0)$

$$R(m_0, \Delta(m_0)) = \frac{2\mu}{\sigma^2}$$
$$\equiv$$
$$-\frac{\mathrm{E}[u''(\cdot)]}{\mathrm{E}[u'(\cdot)]}$$

return to a_0

Exponential Utility Example

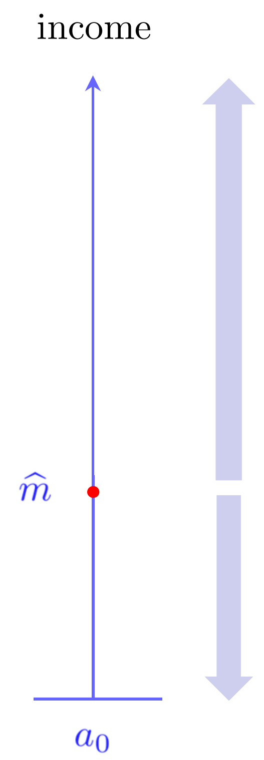

take action $a_0 + \Delta(m_0)$

$$\Delta(m_0) = \frac{2}{\sigma^2}\left(\beta - \frac{2\mu}{\sigma^2}\right)(m_1^* - \widehat{m})$$

return to a_0

$$\widehat{m} = -\frac{1}{\beta}\ln\left(\frac{\alpha}{\beta}\frac{2\mu}{\sigma^2}\left(\beta - \frac{2\mu}{\sigma^2}\right)^{-1}\right)$$

Agenda

- The Model
- Beliefs and Expected Utility
- Managerial Learning
  - The First Period
  - The Second and Subsequent Periods
  - The Optimal Learning Rule
- Persistent Performance Differences
- Barriers to Imitation
- Conclusions

First Period

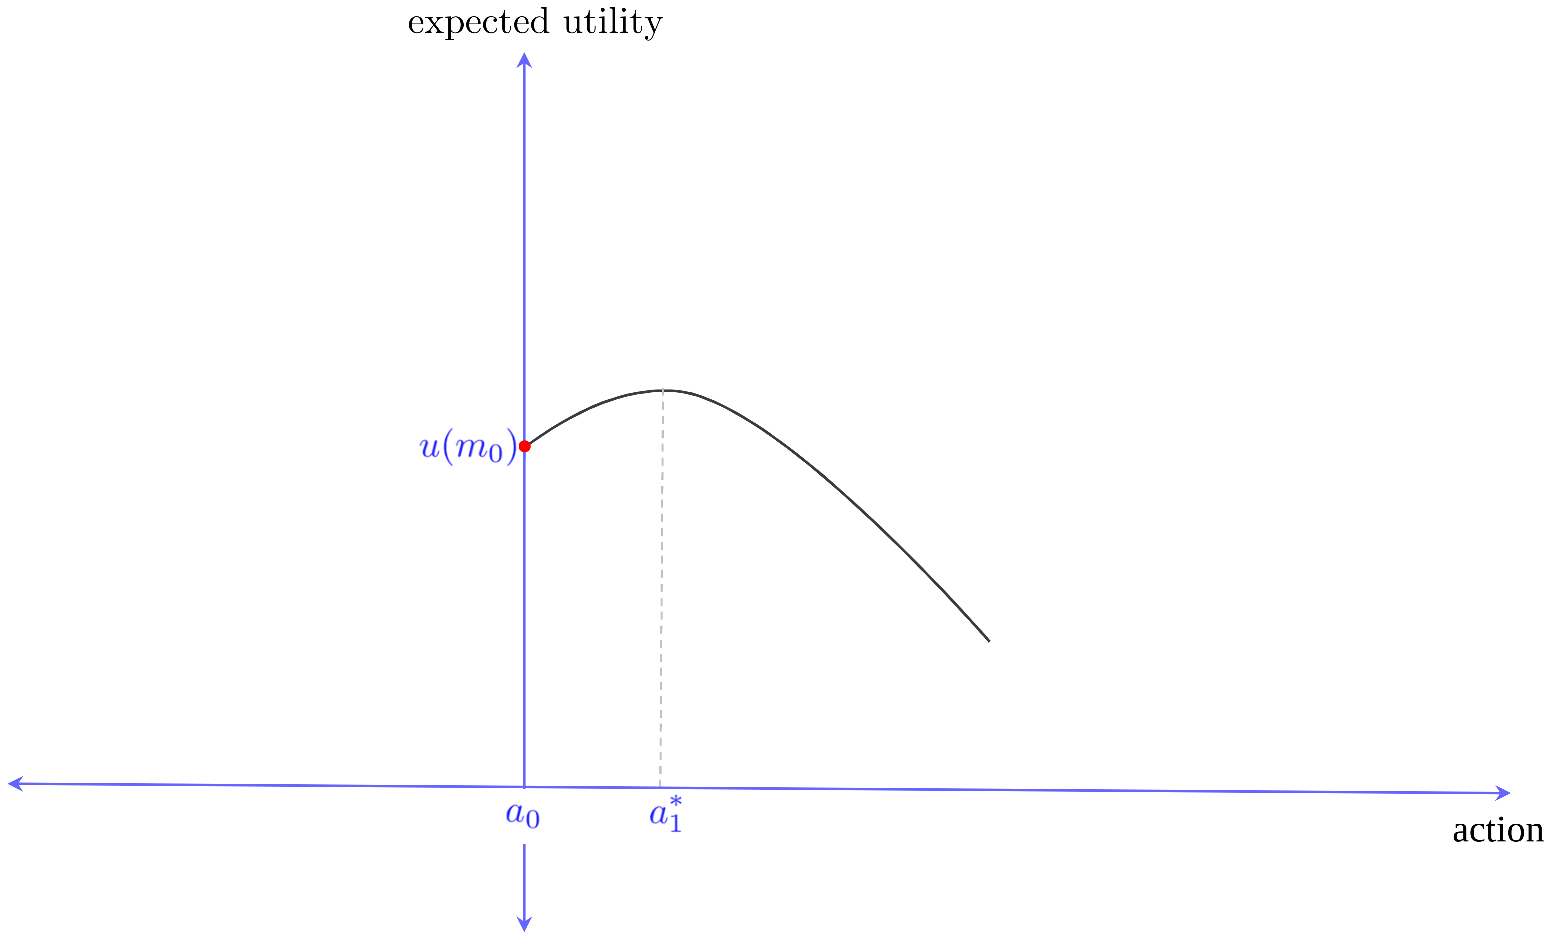

Second Period—Information

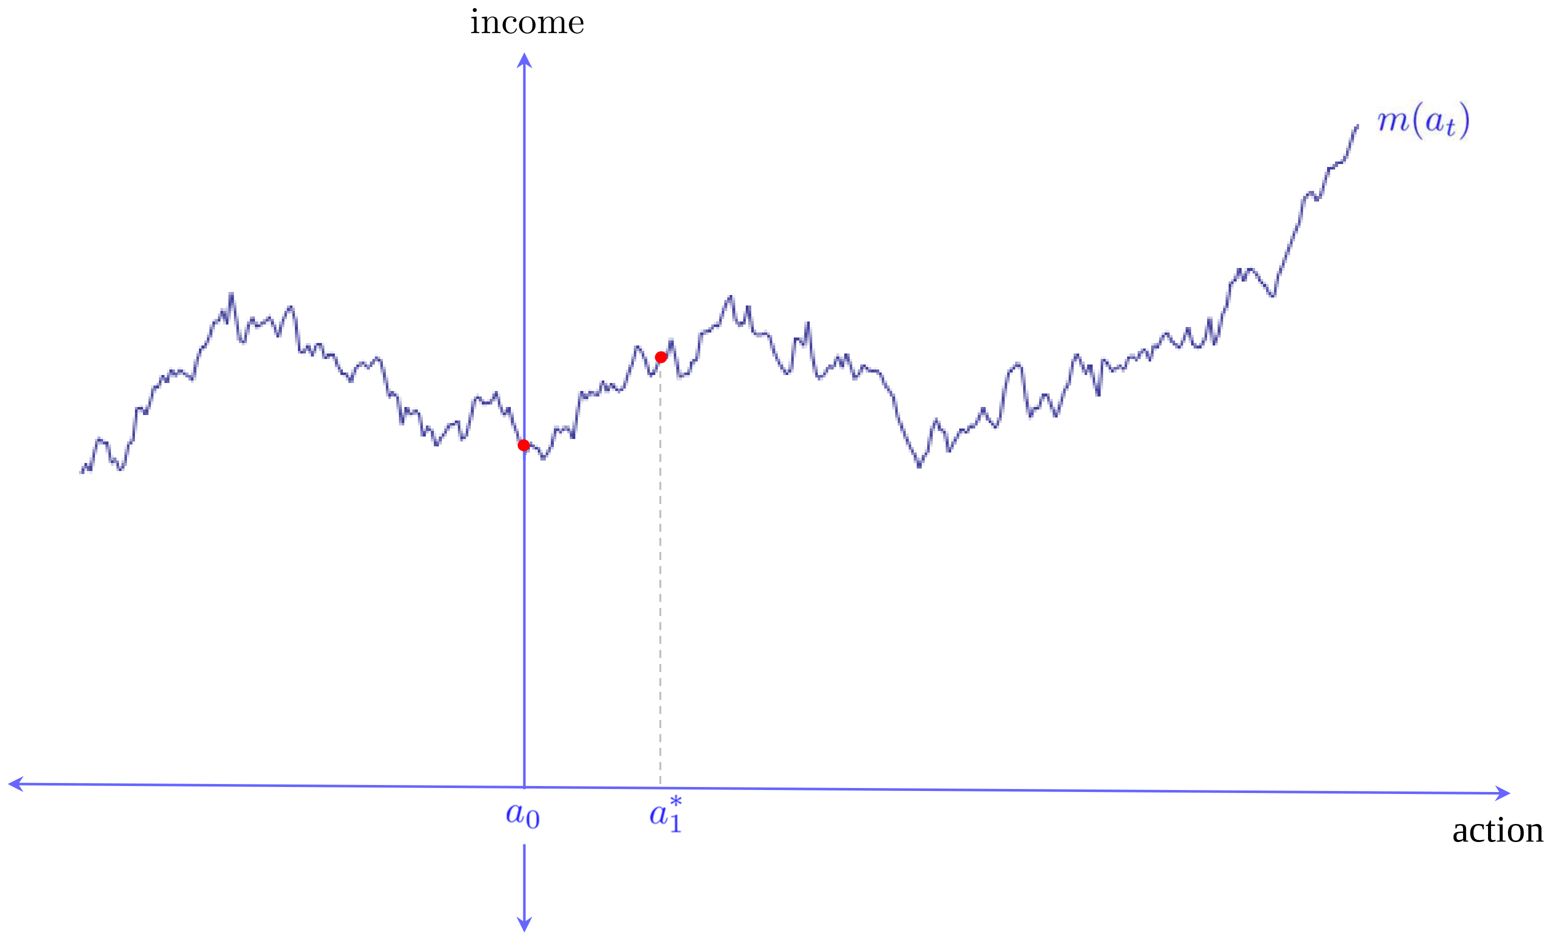

Second Period—Beliefs

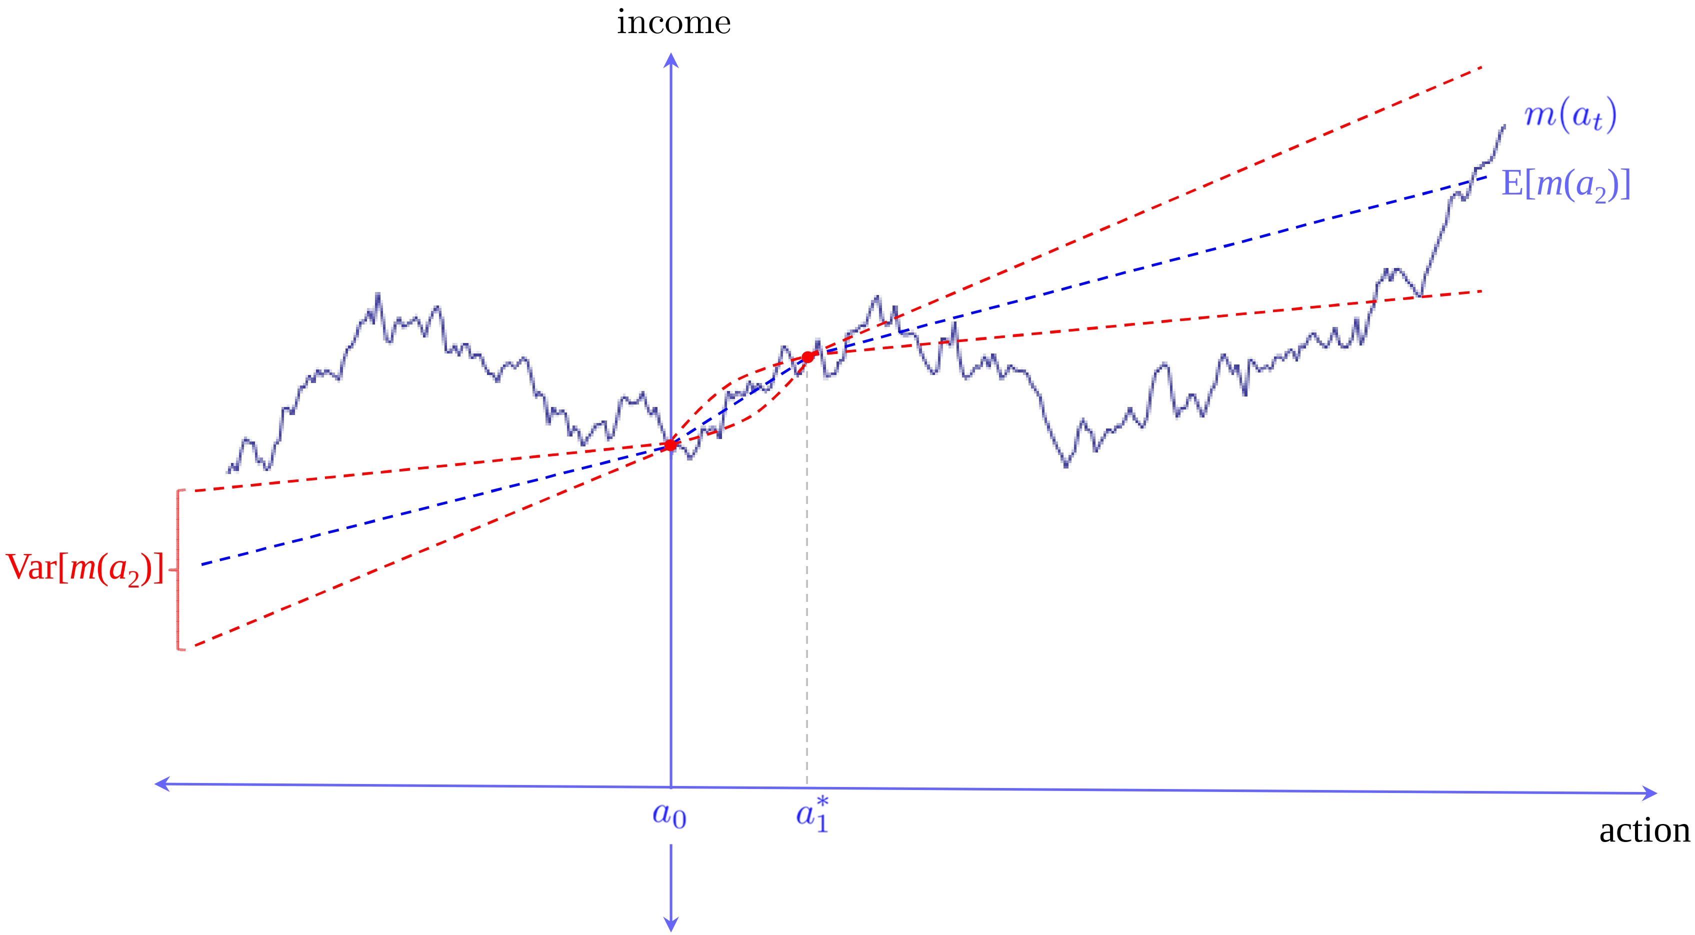

Second Period—The Constrained Problem

$$\max_{\Delta_2 \geq 0} W\left(m_1^* + \mu\Delta_2, \sigma^2\Delta_2\right)$$

First Period—The Problem

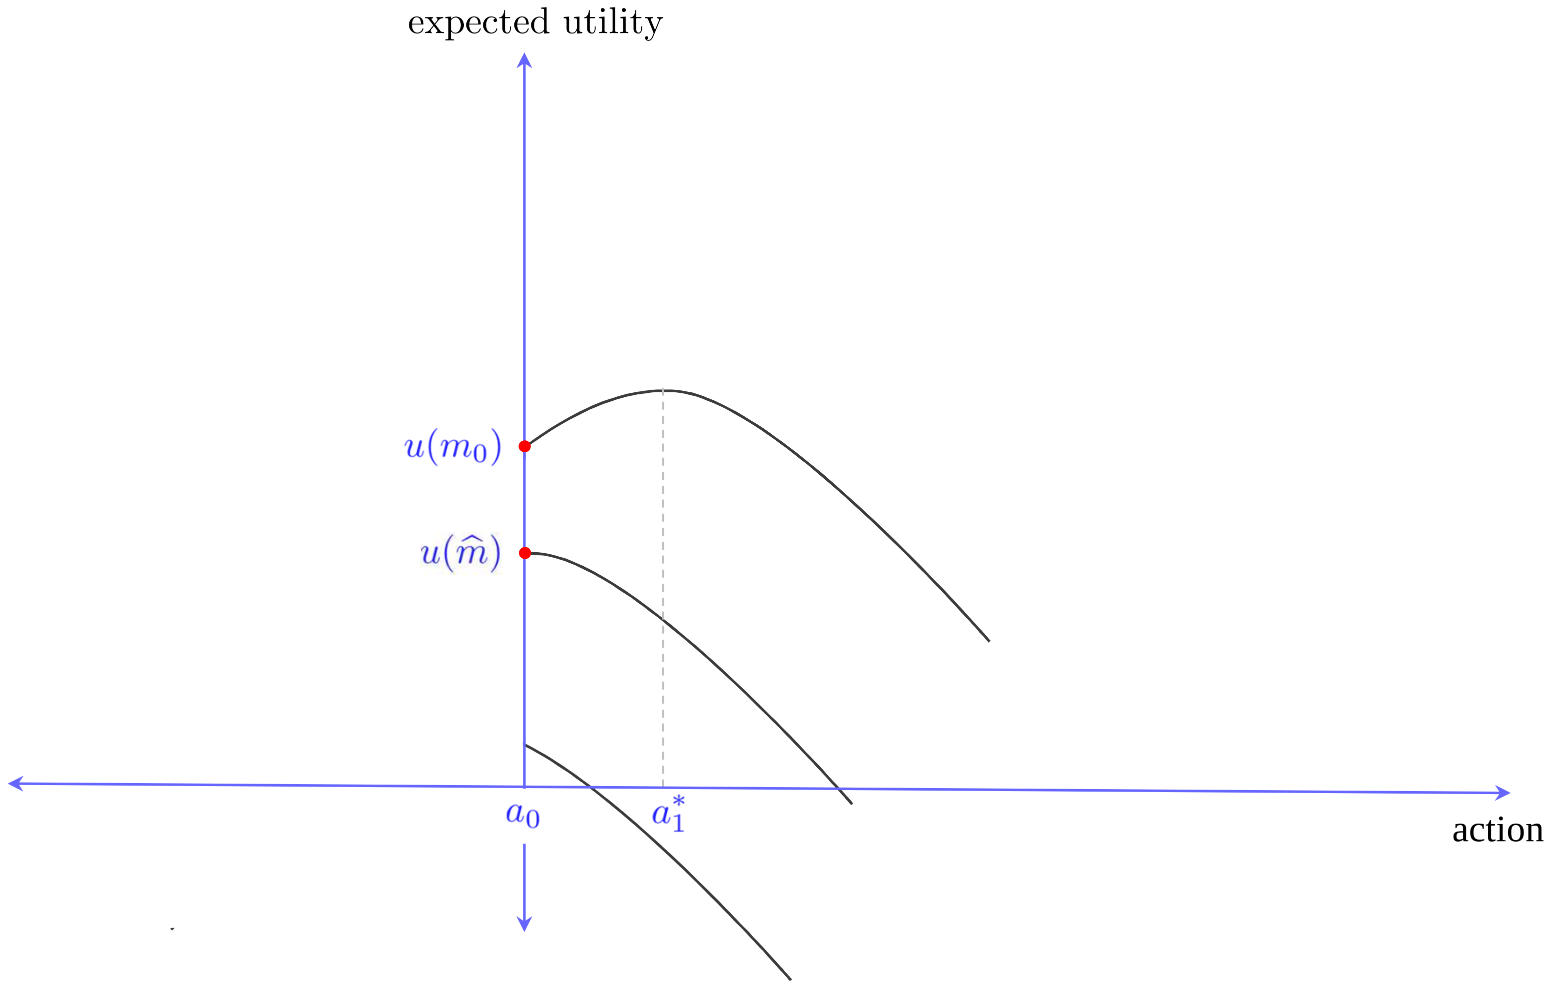

Second Period—The Constrained Problem

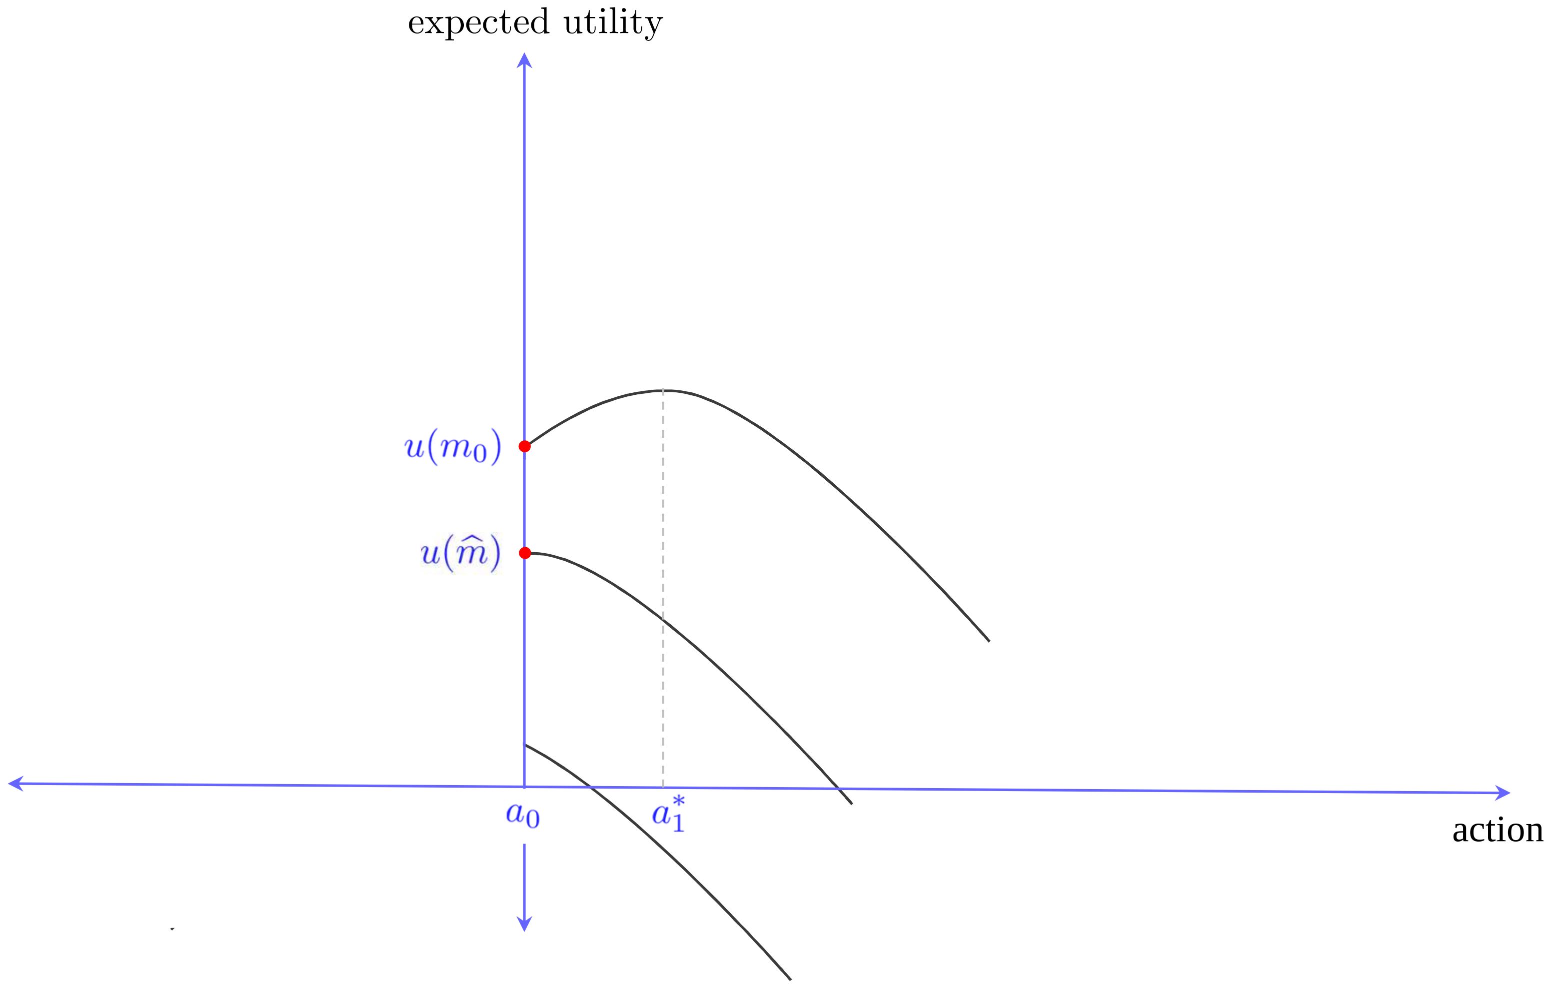

Second Period—The Full Problem

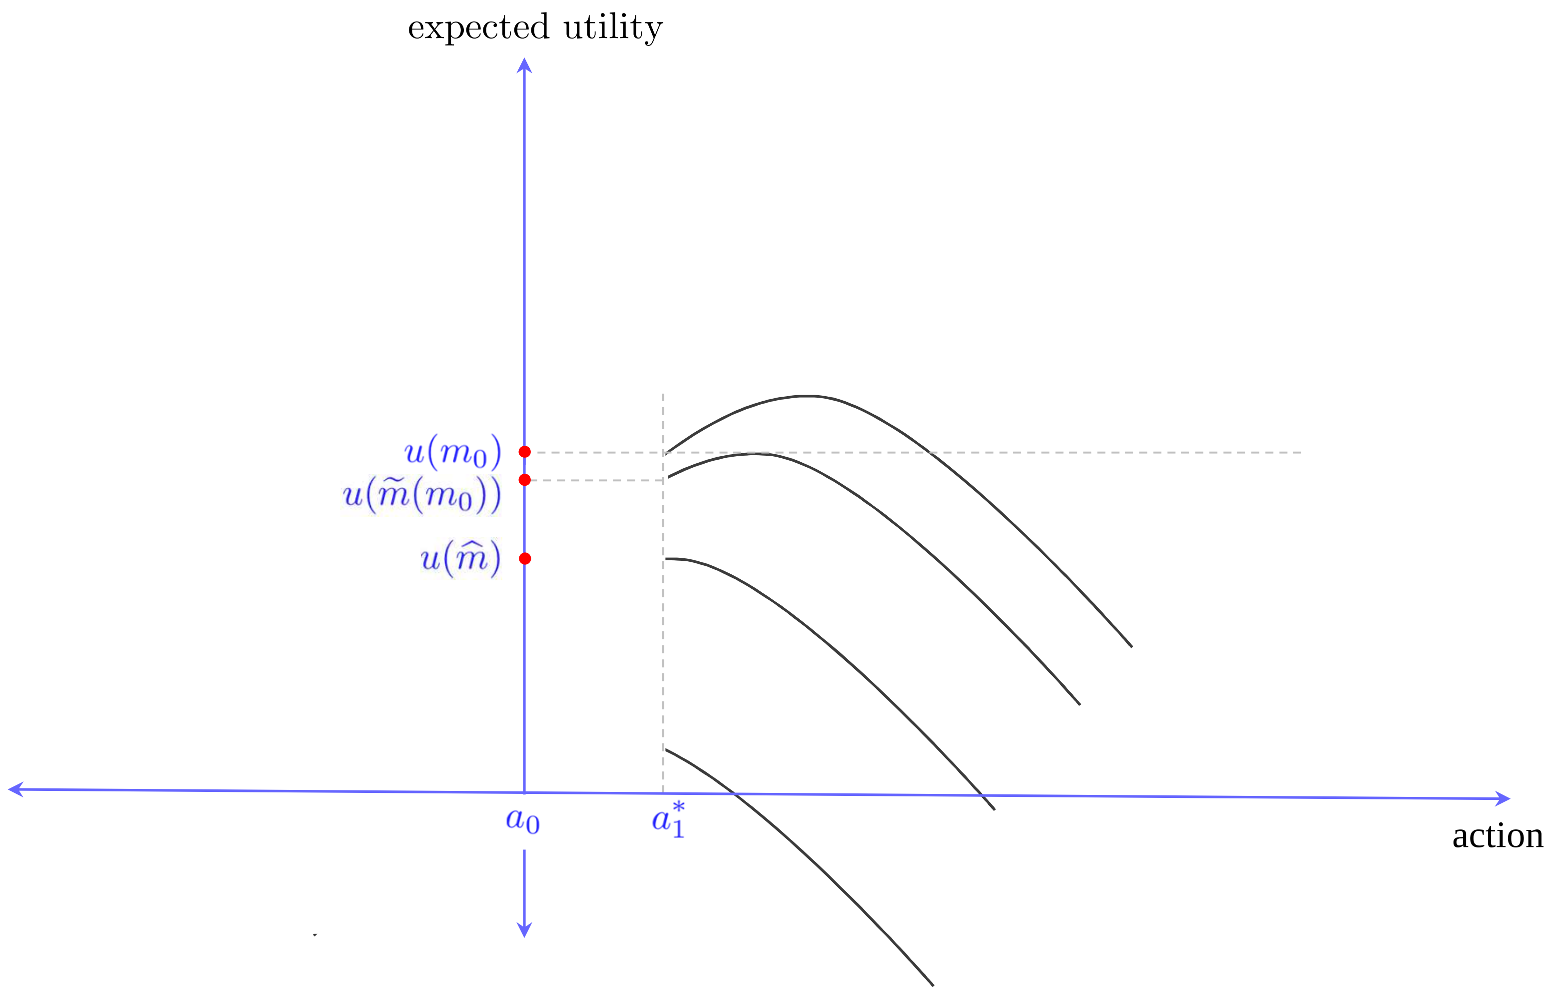

Second Period—The Full Problem

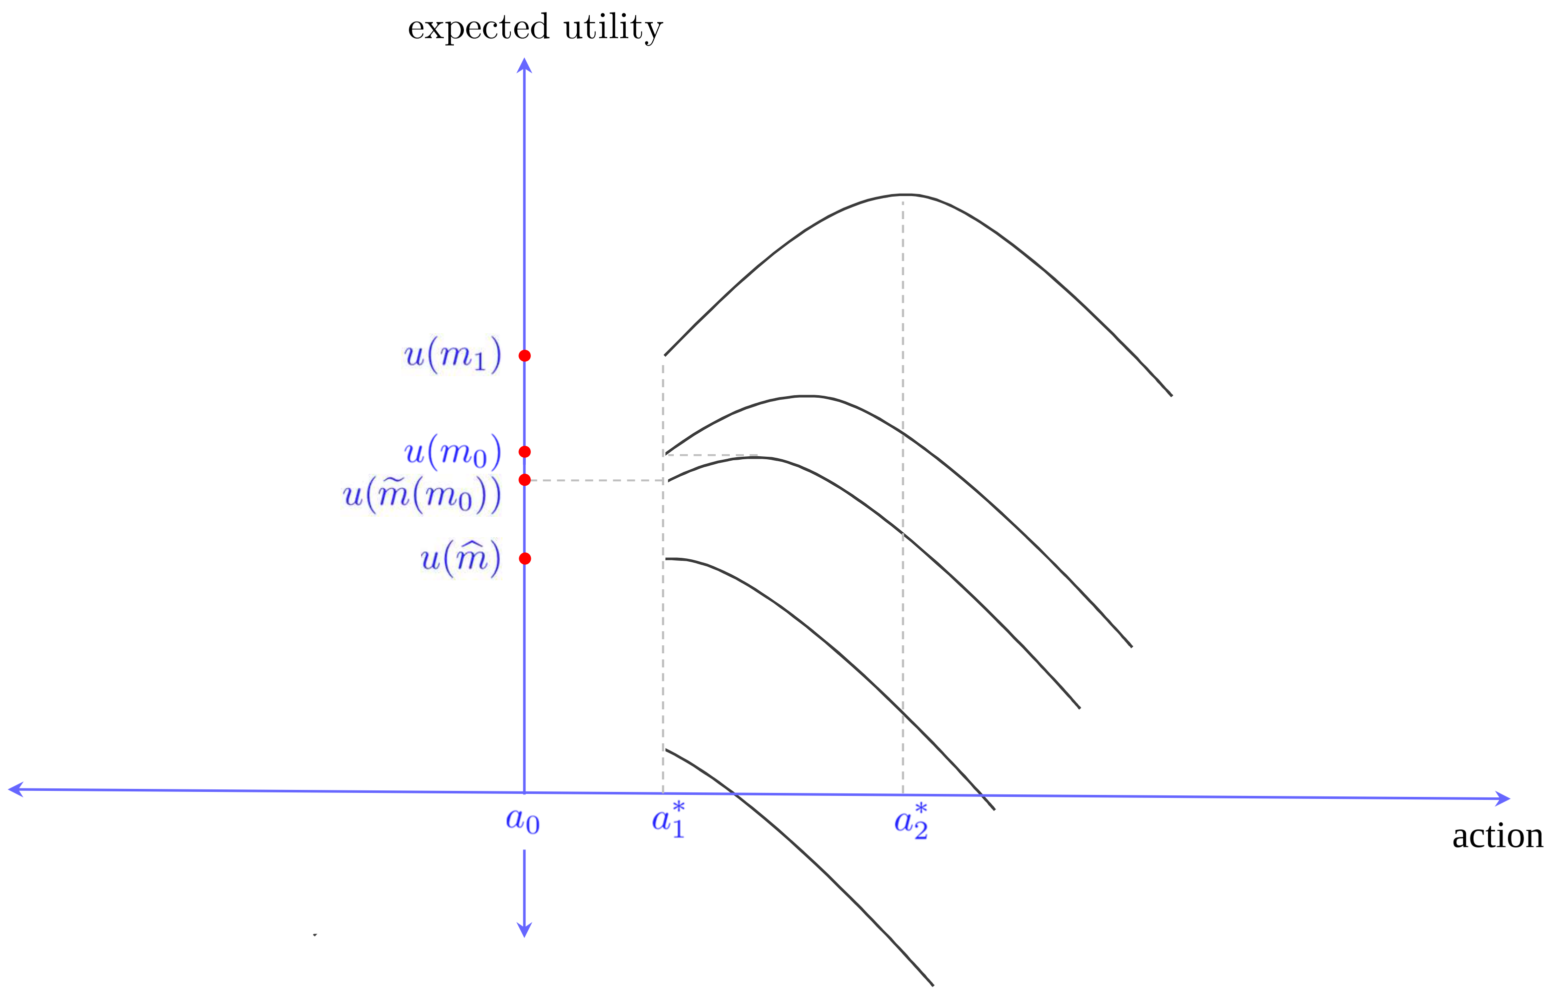

Optimal Learning—First Period

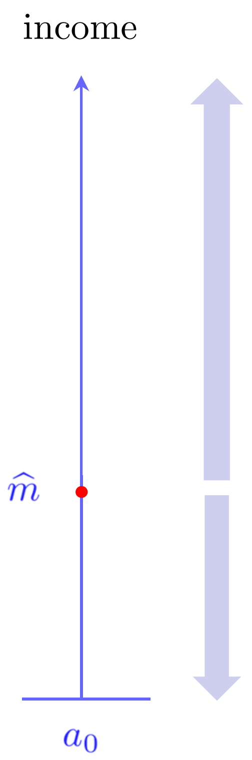

take action $a_0 + \Delta(m_0)$

$$R\left(m_0, \Delta\left(m_0\right)\right) = \frac{2\mu}{\sigma^2}$$

return to a_0

Optimal Learning—Second Period

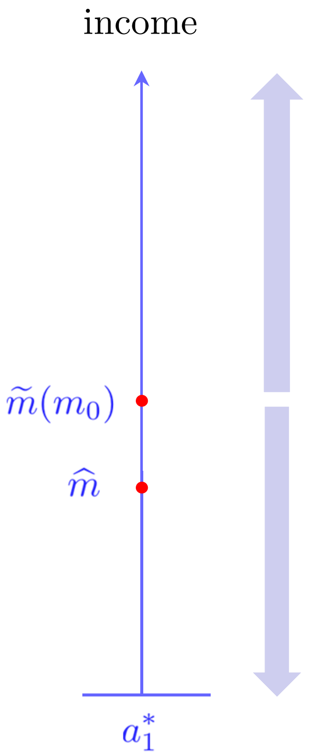

take action $a_1^* + \Delta(m_1^*)$

$$R\left(m_1^*, \Delta\left(m_1^*\right)\right) = \frac{2\mu}{\sigma^2}$$

return to a_0

Exponential Utility Example

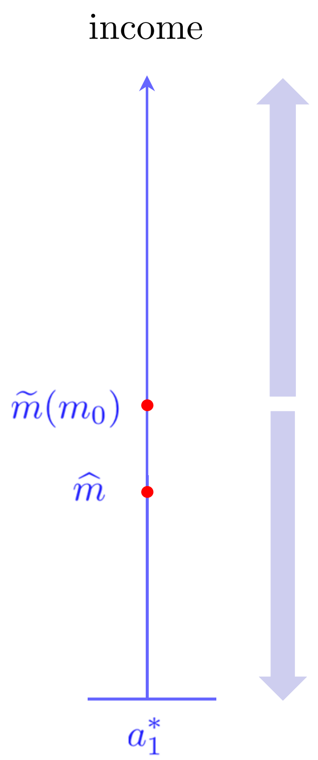

take action $a_1^* + \Delta(m_1^*)$

$$\Delta(m_1^*) = \frac{2}{\sigma^2}\left(\beta - \frac{2\mu}{\sigma^2}\right)(m_1^* - \widehat{m})$$

return to a_0

$$\widetilde{m} = \widehat{m} + (\exp(-\beta\widehat{m}) - \exp(-\beta m_0)) \,/\, \left(\alpha + \alpha\frac{2\mu}{\sigma^2}\left(\beta - \frac{2\mu}{\sigma^2}\right)^{-1}\right)$$

Optimal Learning—Second Period

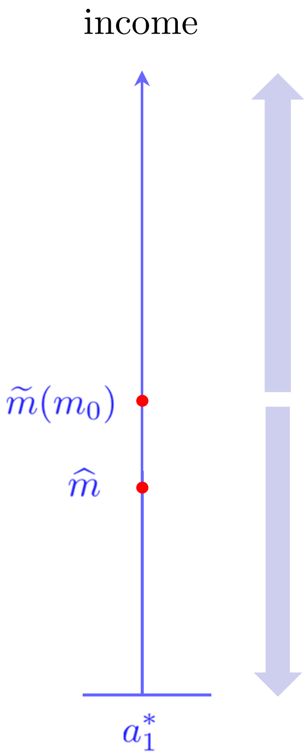

take action $a_1^* + \Delta(m_1^*)$

$$R\left(m_1^*, \Delta\left(m_1^*\right)\right) = \frac{2\mu}{\sigma^2}$$

return to a_0

Optimal Learning—Period t

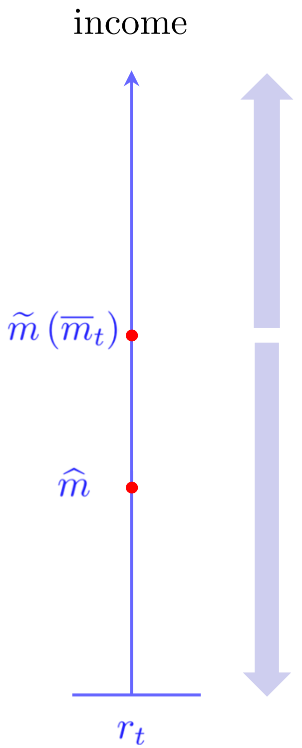

take action $r_t + \Delta\left(m\left(r_t\right)\right)$

$$R\left(m(r_t), \Delta\left(m\left(r_t\right)\right)\right) = \frac{2\mu}{\sigma^2}$$

return to $a\left(\overline{m}_t\right)$

Agenda

- The Model
- Beliefs and Expected Utility
- Managerial Learning
    - The First Period
    - The Second and Subsequent Periods
    - The Optimal Learning Rule
- Persistent Performance Differences
- Barriers to Imitation
- Conclusions

The Optimal Learning Rule

- Managerial search:
  - If starting point is too low, stay at status quo.
  - If starting point is sufficiently high, start marching to the right.
  - Continue marching to right until you fall of sufficiently large cliff.
  - Once that happens, return to the highest known peak.
- Note that:
  - Manager is not satisficing.

The Optimal Learning Rule

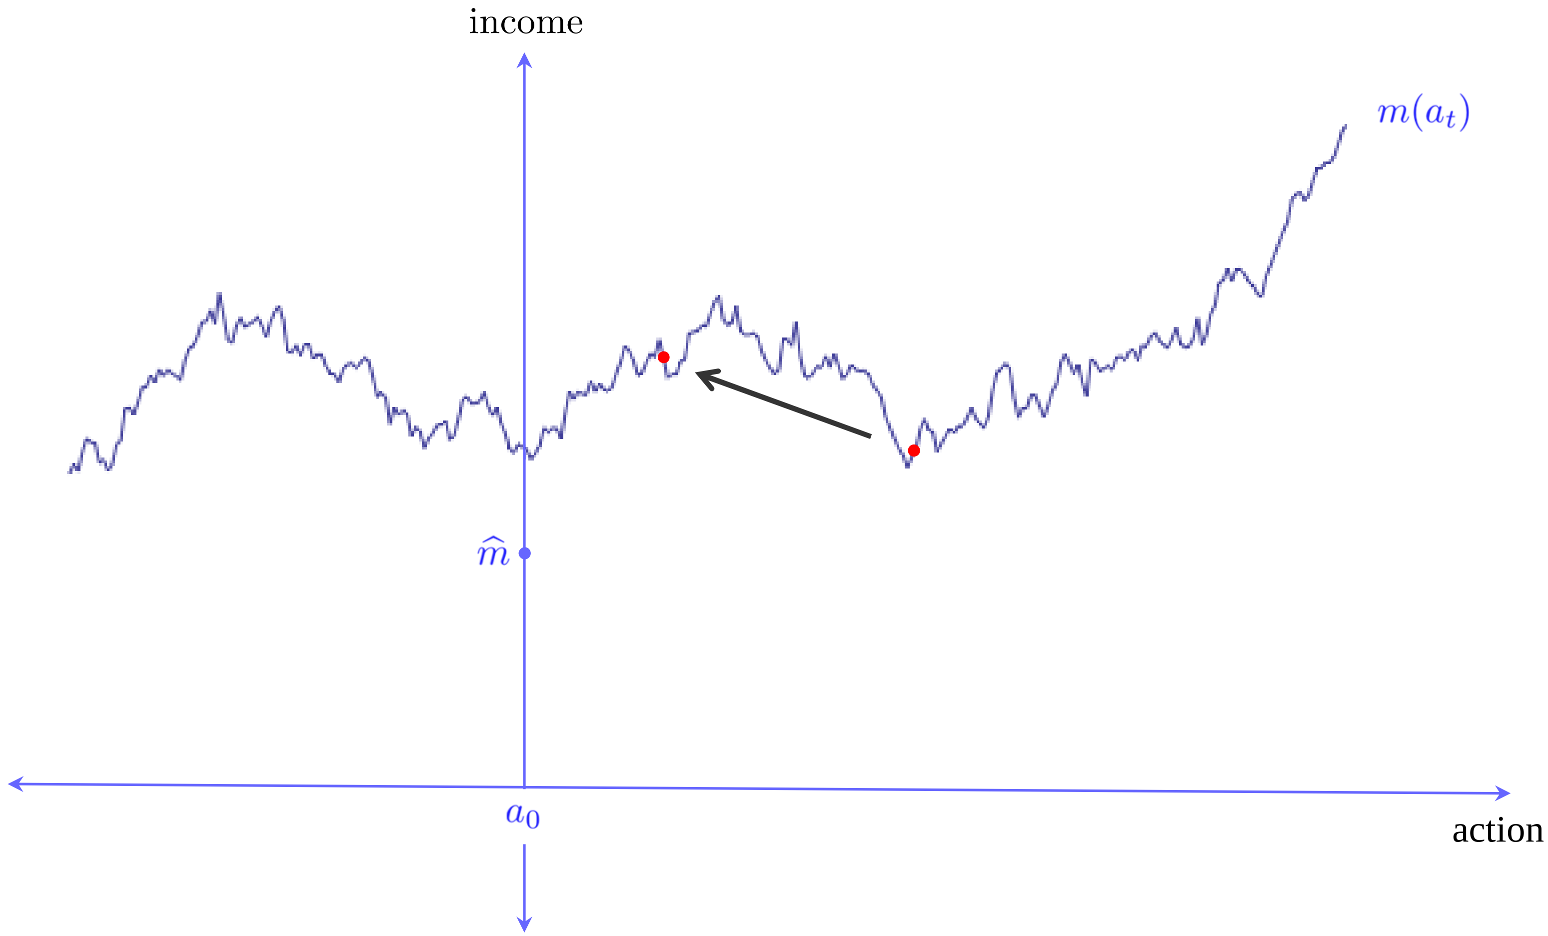

The Optimal Learning Rule

- Managerial search:
  - If starting point is too low, stay at status quo.
  - If starting point is sufficiently high, start marching to the right.
  - Continue marching to right until you fall of sufficiently large cliff.
  - Once that happens, return to the highest known peak.
- Note that:
  - Manager is not satisficing.
  - In general, manager does not end up at local peak.

The Optimal Learning Rule

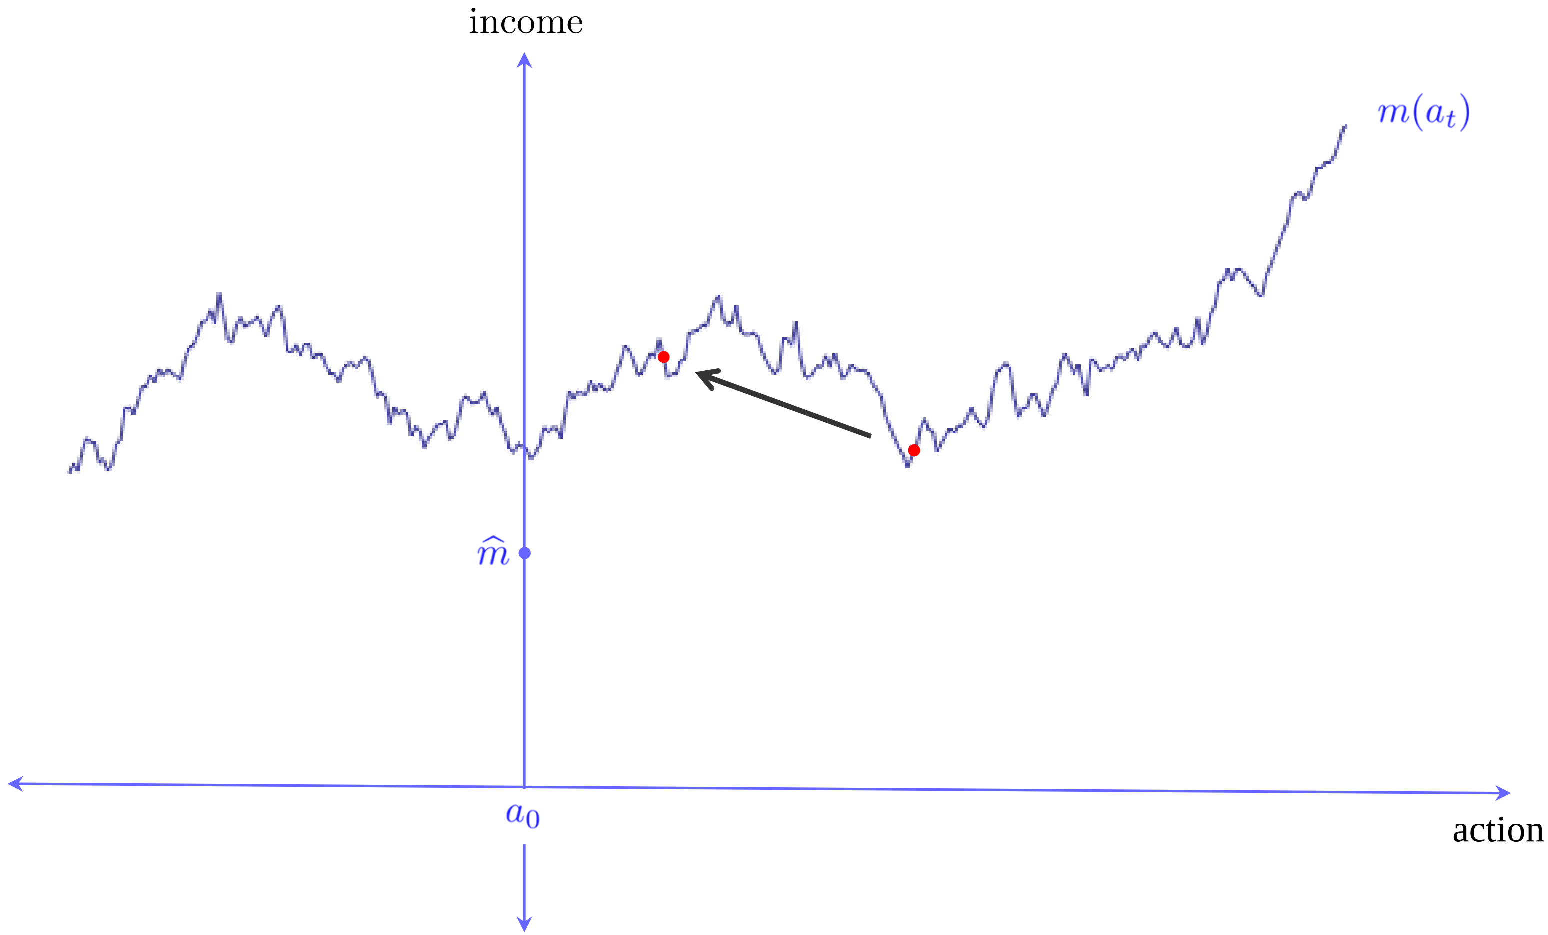

The Optimal Learning Rule

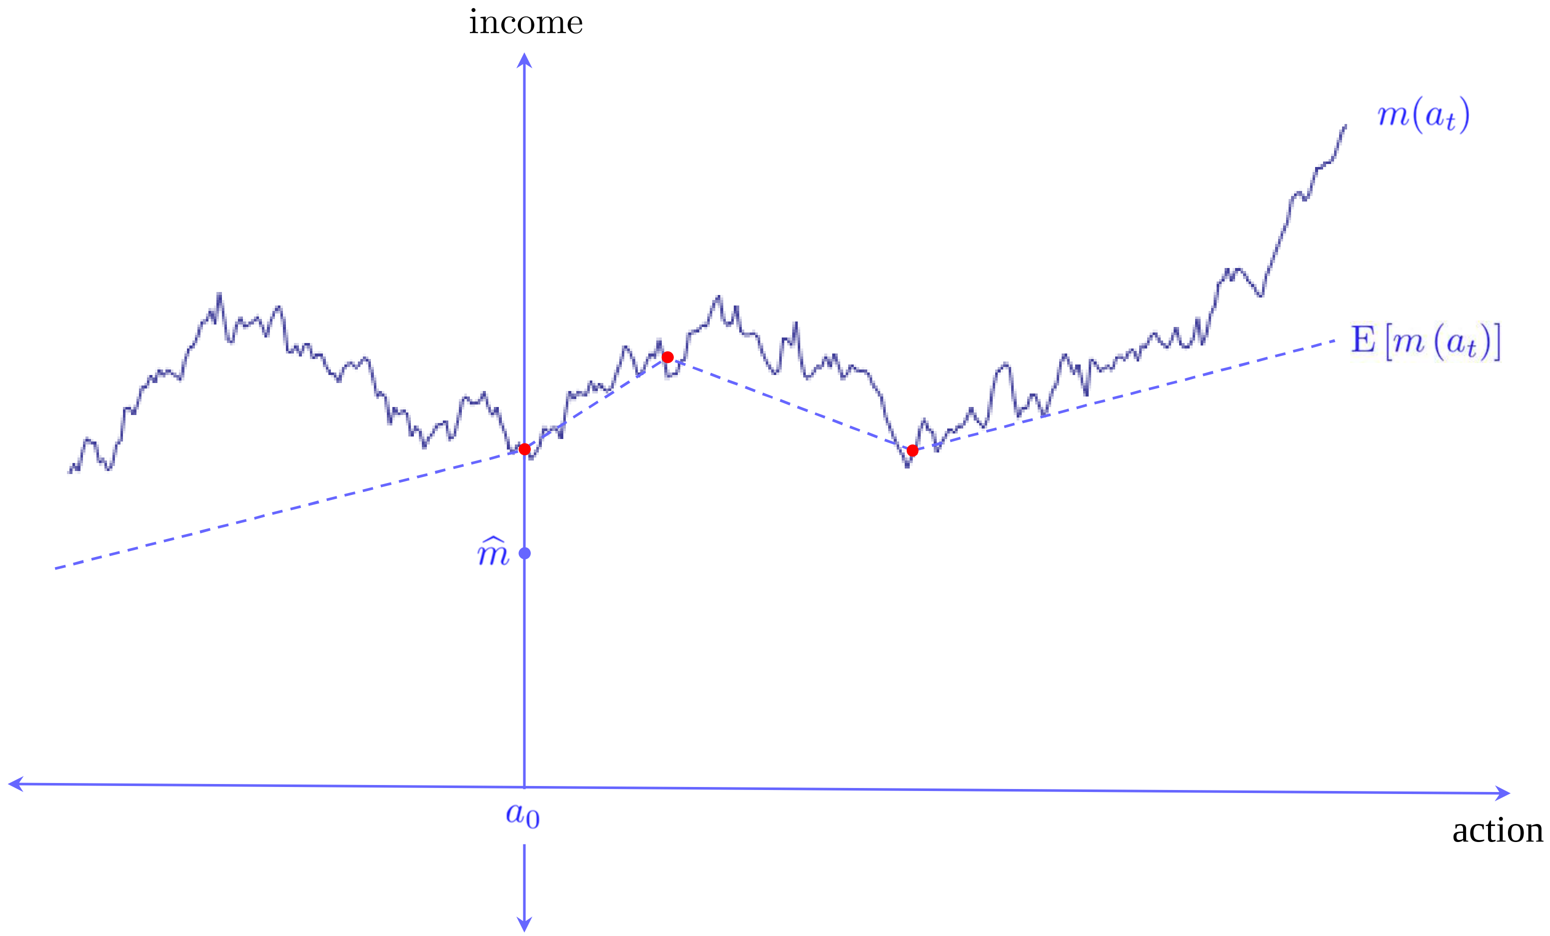

Agenda

- The Model
- Beliefs and Expected Utility
- Managerial Learning
- Persistent Performance Differences
    - Different Initial Conditions
    - Complementarities and Decentralized Learning
- Barriers to Imitation
- Conclusions

Different Initial Conditions

- Managerial search depends crucially on status quo income.
- Rich dynamics are possible: convergence, divergence, overtaking...

Example of Overtaking

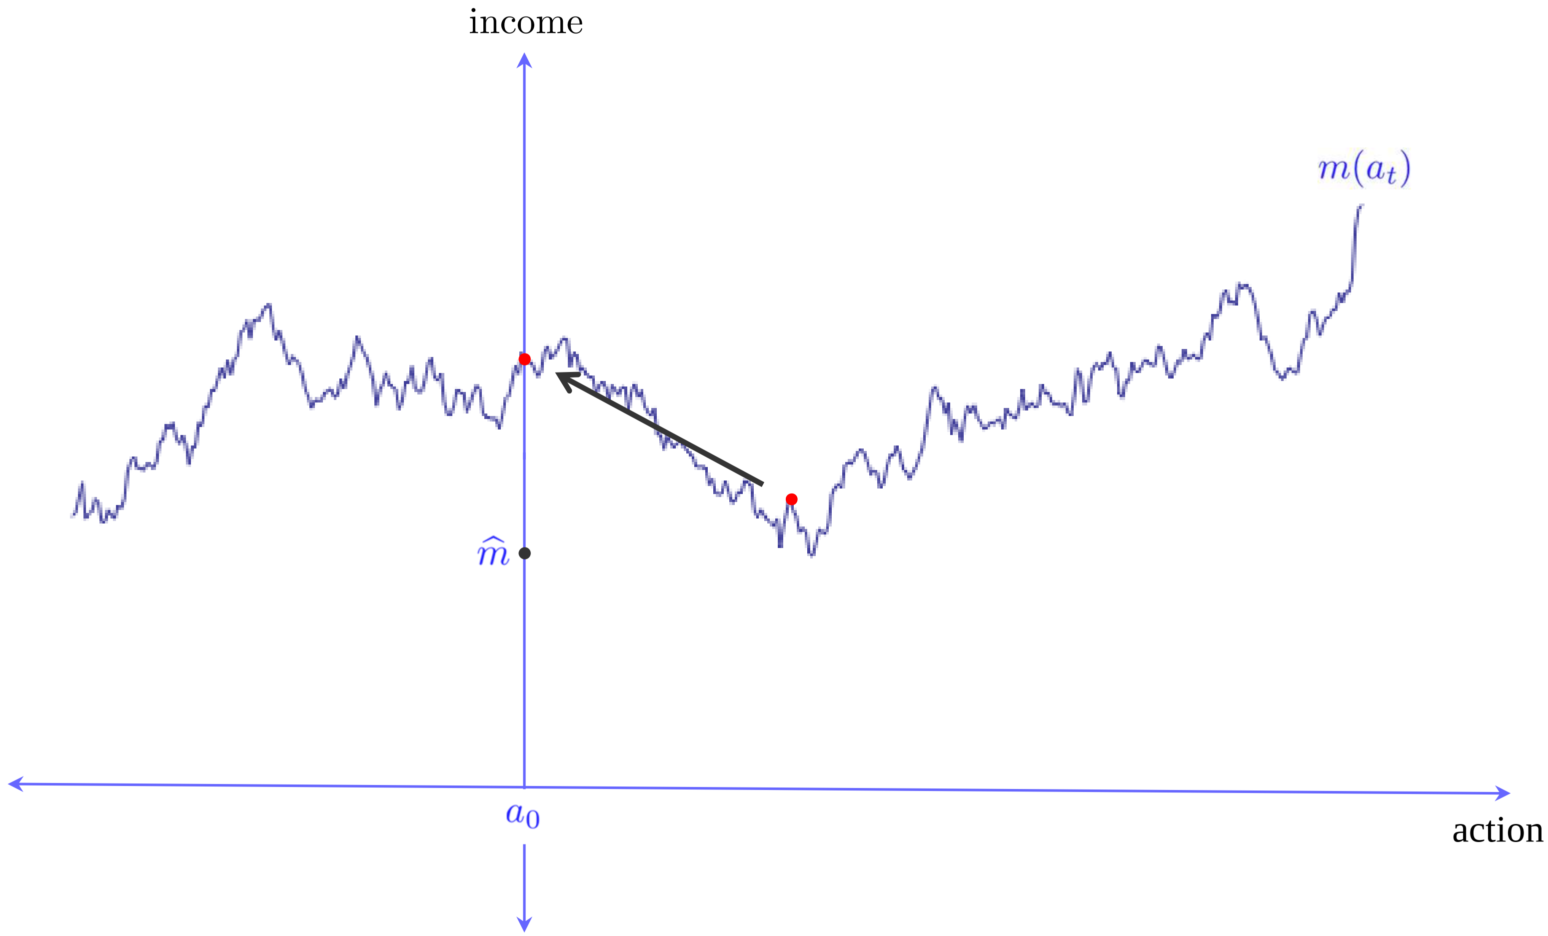

Example of Overtaking

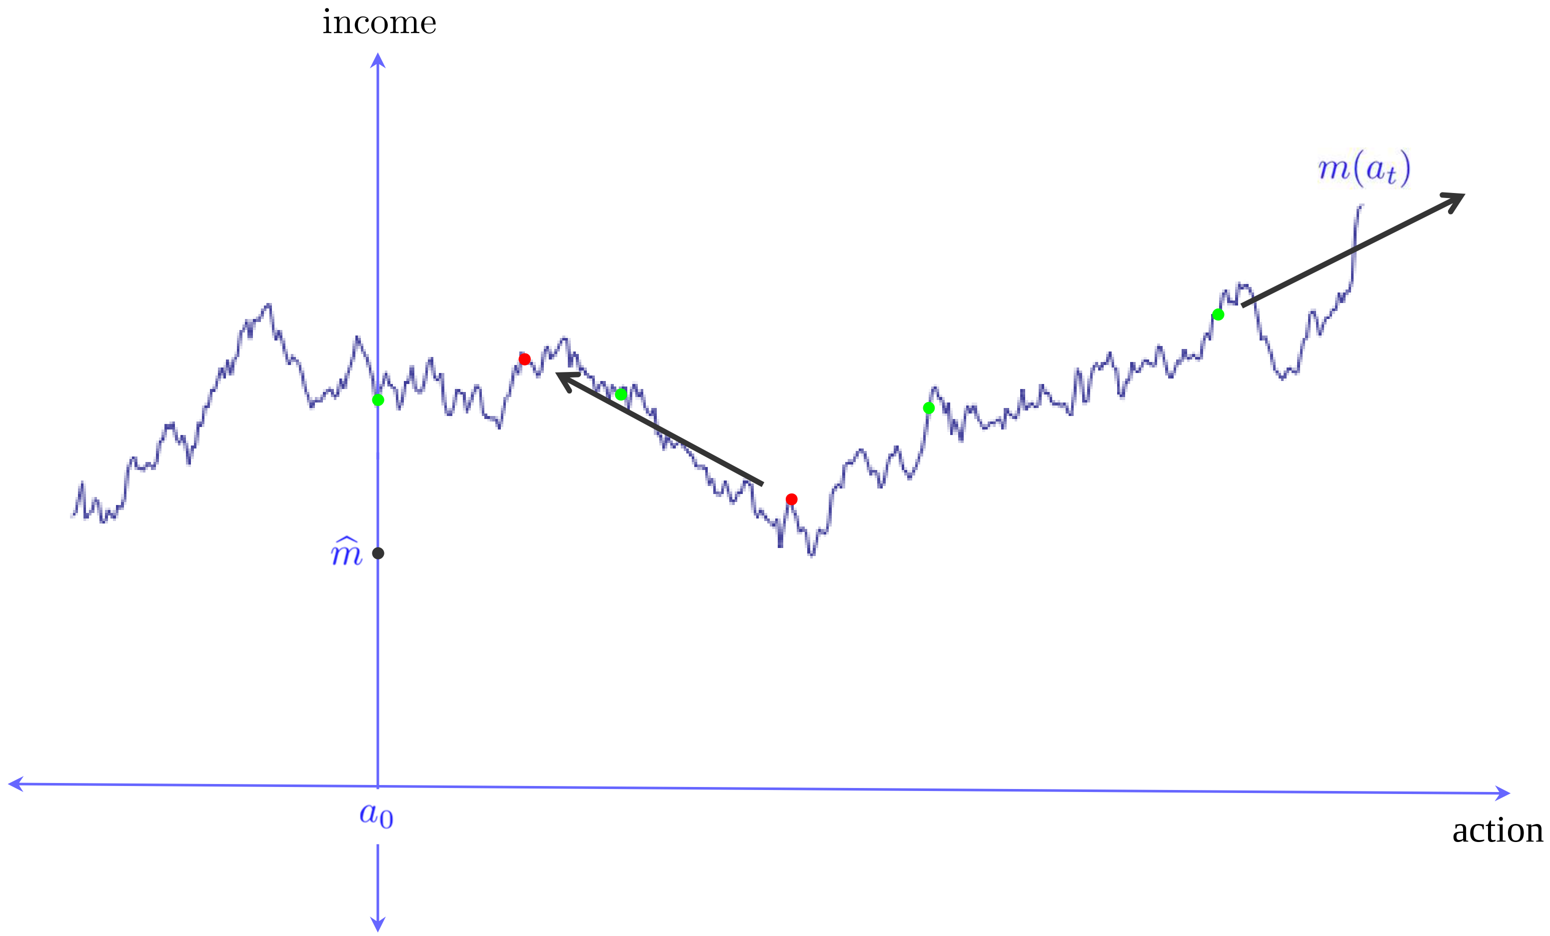

Different Initial Conditions

- An increase in status quo income...

...makes it more likely that the manager engages in search in any period t:

$$\frac{\mathrm{d}prob\left(m_t^* > \widetilde{m}\left(\overline{m}_t\right)\right)}{\mathrm{d}m_0} > 0$$

...leads to an even larger increase in expected income in any period t:

$$\frac{\mathrm{dE}\left[m_t^*(m_0)\right]}{\mathrm{d}m_0} > 1$$

→ Differences in status quo incomes do not only persist, they actually grow larger.

Agenda

- The Model
- Beliefs and Expected Utility
- Managerial Learning
- Persistent Performance Differences
  - Different Initial Conditions
  - Complementarities and Decentralized Learning
- Barriers to Imitation
- Conclusions

Complementarities and Decentr. Learning

- Two managers, A and B.
- Each manager's income is given by

degree of complementarity

$$m_t^A = m\left(a_t^A\right) - \frac{1}{2}\delta\left(a_t^A - a_t^B\right)^2$$

$$m_t^B = m\left(a_t^B\right) - \frac{1}{2}\delta\left(a_t^A - a_t^B\right)^2$$

- The managers know the properties of the Brownian motion.
- And they know the income associated with any action either of them took in the past.

Optimal Learning—First Period

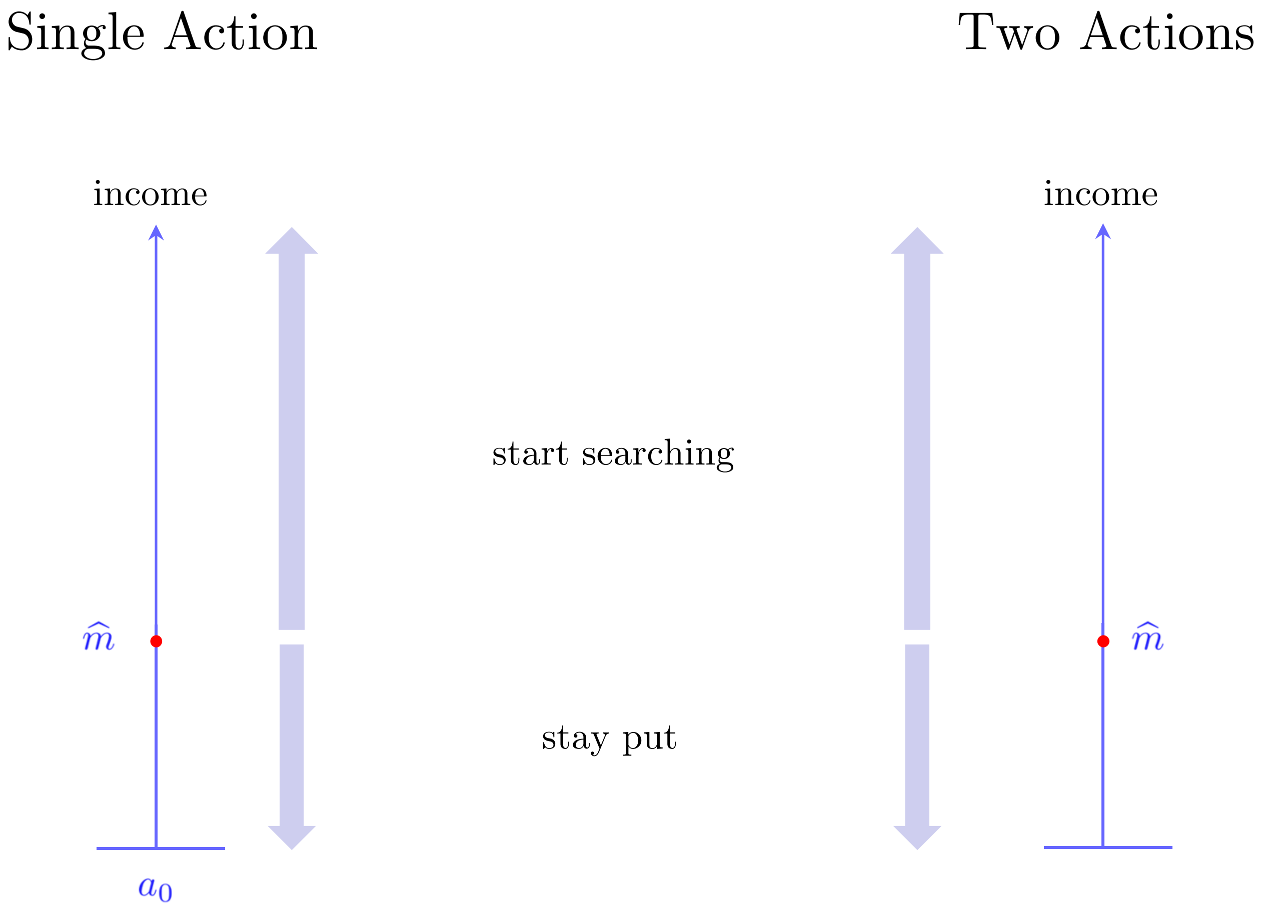

Optimal Learning—Second Period

Single Action | Two Actions

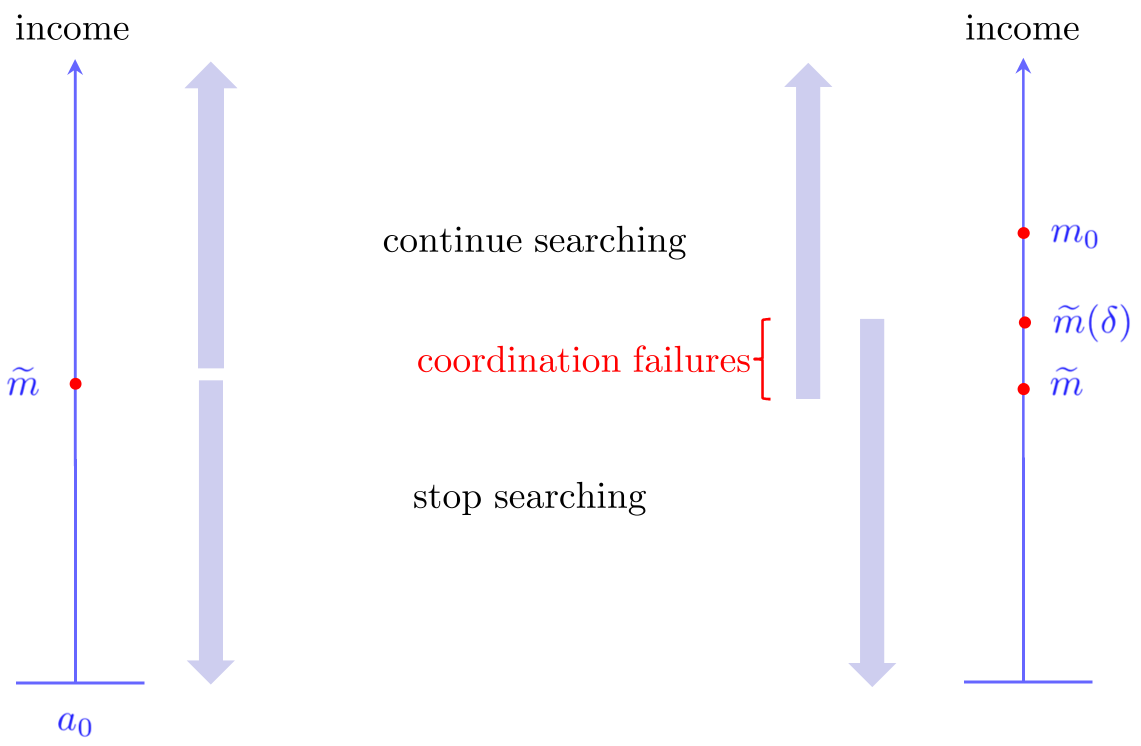

Complementarities and Decentr. Learning

- First period learning the same as in single action model.
- But in later periods coordination failures can arise.
- Coordination failures arise for intermediate income levels.
- Coordination failures are more likely, the stronger the compl.
- If actions are very complementary, the only way to avoid coordination failures is for income to go up every period.
- Implications for "big push" policies.

Agenda

- The Model
- Beliefs and Expected Utility
- Managerial Learning
- Persistent Performance Differences
- Barriers to Imitation
- Conclusions

Barriers to Imitation

- One manager.
- Two actions: a_t^A and a_t^B
- Actions are strict complements: $a_t^A = a_t^B = a_t$
- Income: $m_t = m^A(a_t) + m^B(a_t)$

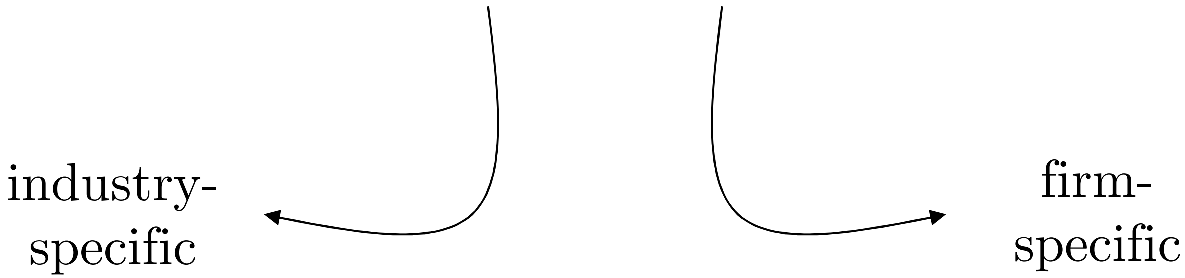

Barriers to Imitation

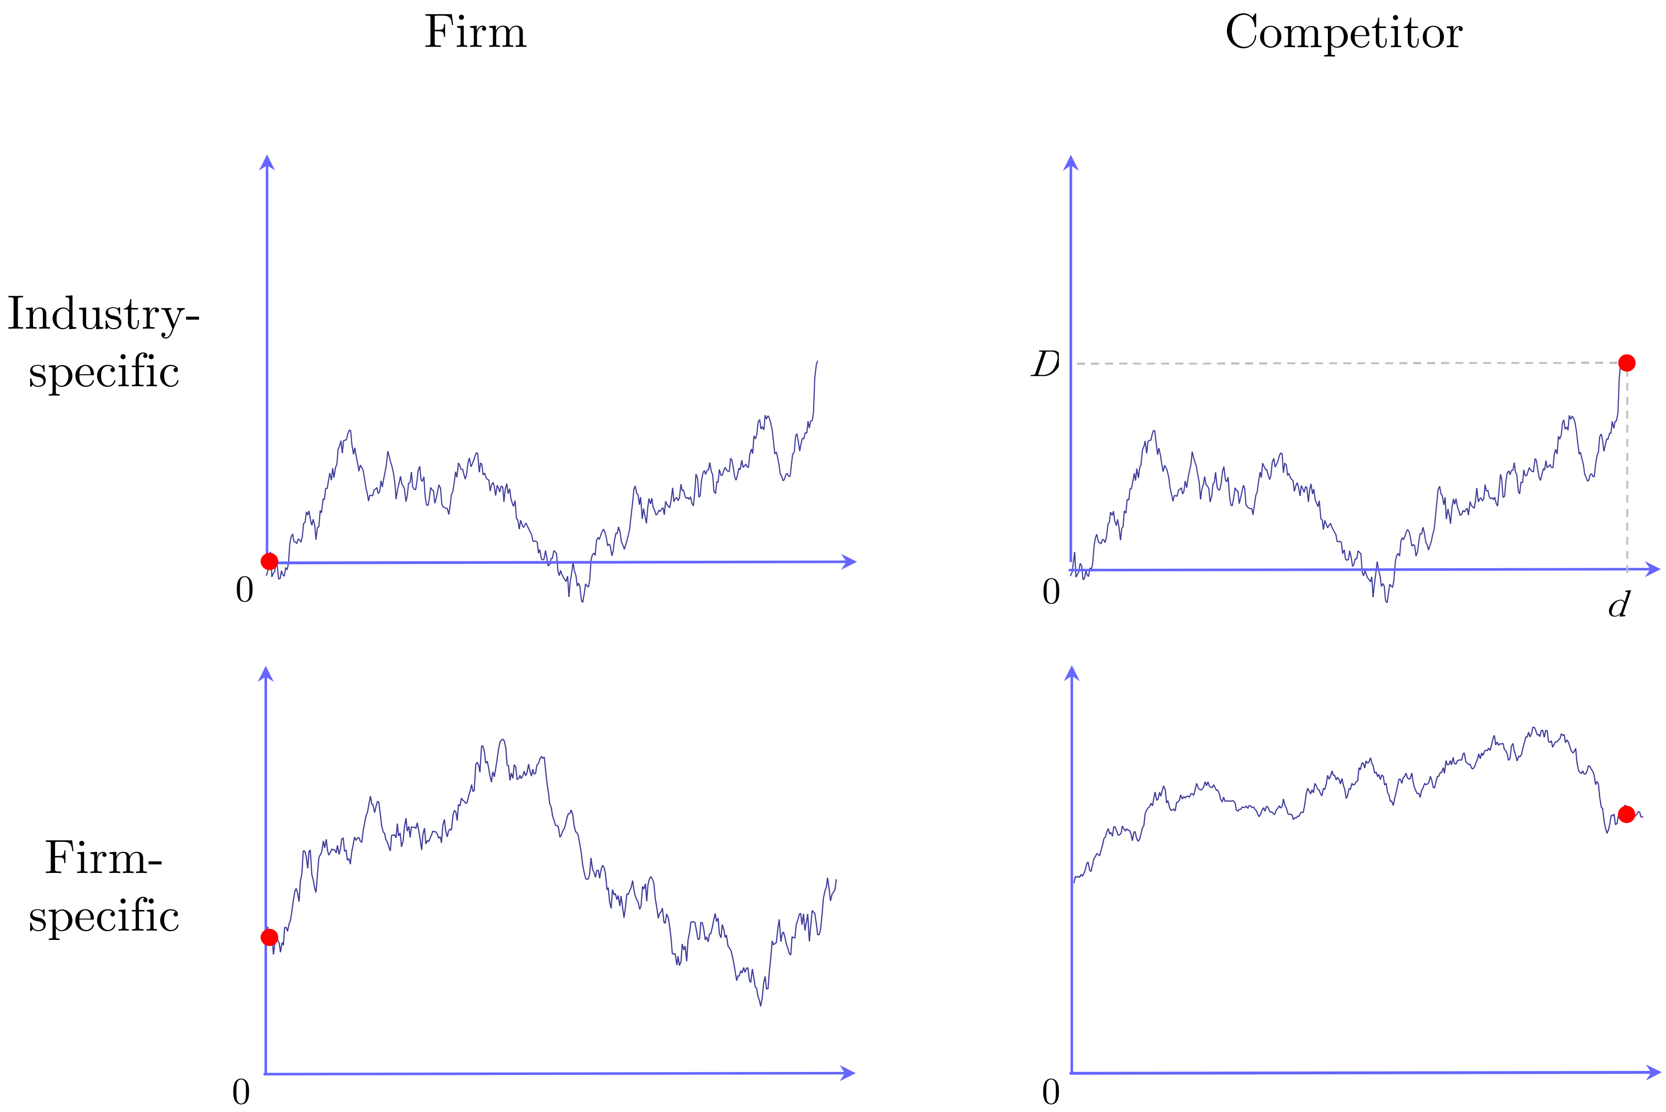

Barriers to Imitation

The manager prefers status quo to imitation if and only if

$$d \geq \widehat{d},$$

where $\widehat{d}$ is increasing in D and decreasing in σ^2.

Agenda

- The Model
- Beliefs and Expected Utility
- Managerial Learning
- Persistent Performance Differences
- Barriers to Imitation
- Conclusions

Conclusions

- Our premise:
  - ➢ Managerial problem too complex to be solved analytically.
  - ➢ Instead, managers learn by trial-and-error.
- Our main results:
  1. Optimal learning rule:
     - ➢ Takes a simple form that generates rich dynamics.
  2. Persistent performance differences:
     - ➢ Can arise even if firms are identical.
  3. Barriers to imitation:
     - ➢ Can be too risky.
- Application to growth and development of nations