Cells in the face of uncertainty

Olivier Rivoire (Part 1)

David Lacoste (Part 2)

David Tourigny

INRAE Département MathNum EURIP LEARNING PLANET INSTITUTE

Introduction: bacterial persistence

Viable bacteria after an antibiotic treatment:

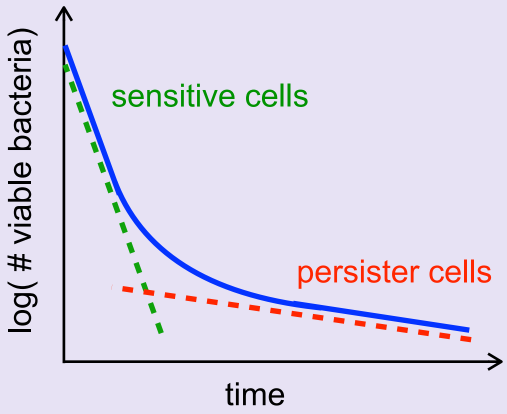

Introduction: bacterial persistence

Viable bacteria after an antibiotic treatment:

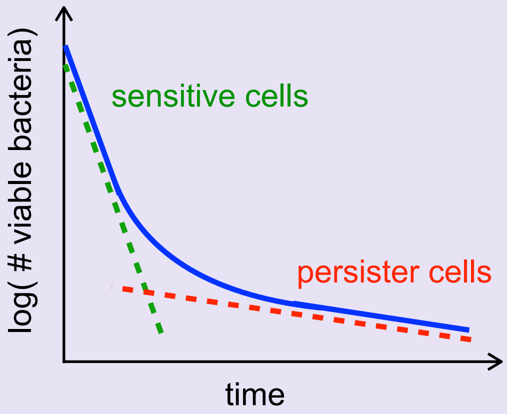

Origin of persisters? - mutations

Introduction: bacterial persistence

Viable bacteria after an antibiotic treatment:

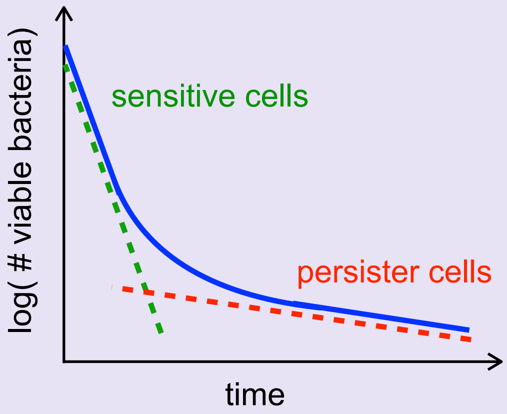

Origin of persisters?
- mutations
- sensing

Introduction: bacterial persistence

Viable bacteria after an antibiotic treatment:

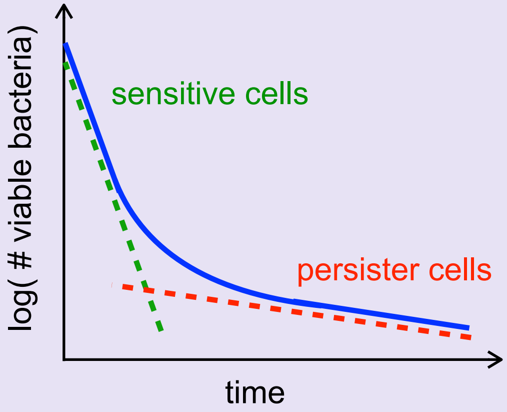

Origin of persisters?
- mutations
- sensing
- phenotypic switching

Introduction: bacterial persistence

Viable bacteria after an antibiotic treatment:

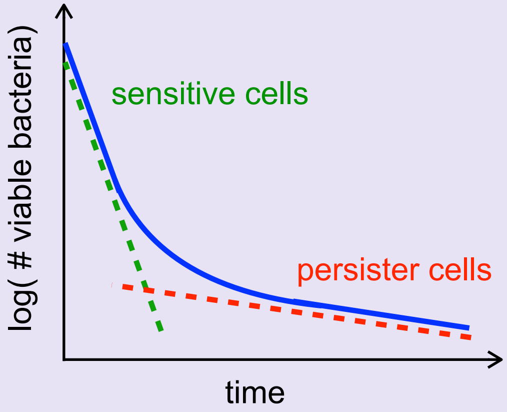

Origin of persisters?
- mutations
- sensing
- phenotypic switching

Phenotypic switching:
- 2 states: growing/sensitive versus dormant/resistant
- same genotype
- both states are present in any environment
- stochastic switches between states

Introduction: bacterial persistence

Viable bacteria after an antibiotic treatment:

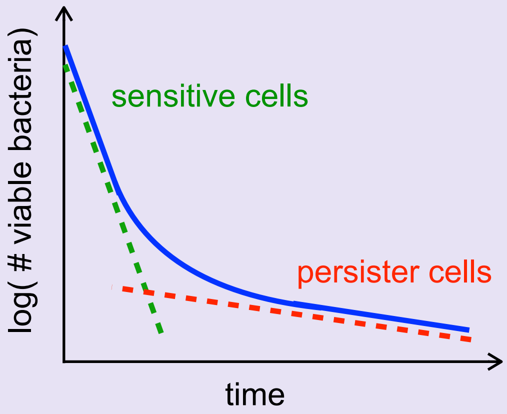

Origin of persisters?

- mutations
- sensing
- phenotypic switching

Phenotypic switching:

- 2 states: growing/sensitive versus dormant/resistant
- same genotype
- both states are present in any environment
- stochastic switches between states

Financial analogy: bet-hedging / portfolio diversification

Introduction: bacterial persistence

Viable bacteria after an antibiotic treatment:

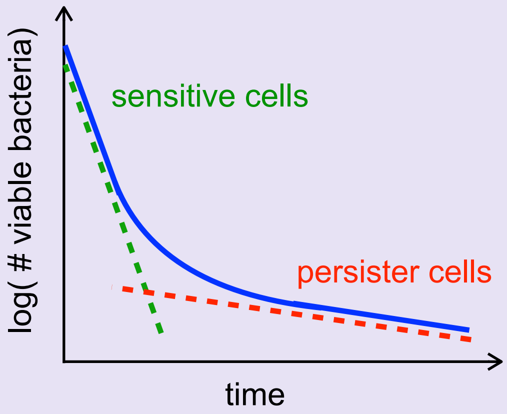

Origin of persisters?
- mutations
- sensing
- phenotypic switching

Phenotypic switching:
- 2 states: growing/sensitive versus dormant/resistant
- same genotype
- both states are present in any environment
- stochastic switches between states

Financial analogy: bet-hedging / portfolio diversification

Question: optimal "strategy"?

Introduction: bacterial persistence

Viable bacteria after an antibiotic treatment:

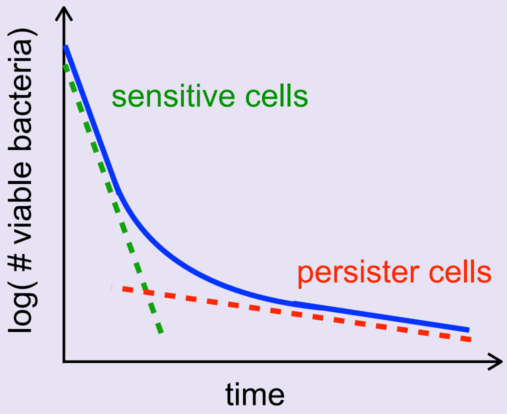

Origin of persisters?
- mutations
- sensing
- phenotypic switching

Phenotypic switching:
- 2 states: growing/sensitive versus dormant/resistant
- same genotype
- both states are present in any environment
- stochastic switches between states

Financial analogy: bet-hedging / portfolio diversification

Question: optimal "strategy"?

Key points:
- optimality in stochastic environments
- individual versus population-level adaptation
- analogies with finance and their limitations

Bacterial persistence: elementary model

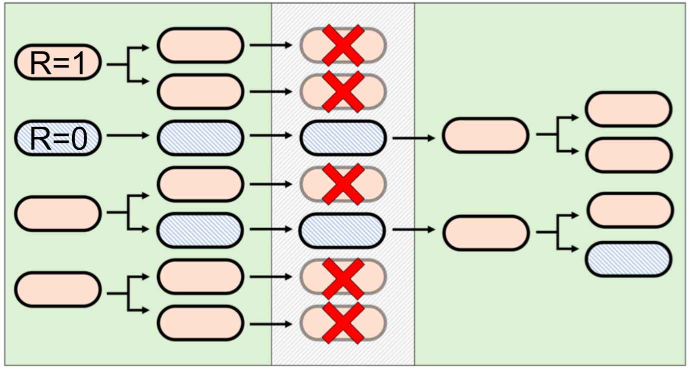

Model assumptions:

- 2 states R: growing (R=1) / dormant (R=0)
- 2 environments E: antibiotic (E=+) / no antibiotic (E=-)
- survival/reproduction per generation f(R,E)

	E = +	E = -
R = 0	1	1
R = 1	0	2

- probability for antibiotic (E=+) : *p*
- probability to be dormant (R=0): *u*

(per generation)

Bacterial persistence: elementary model

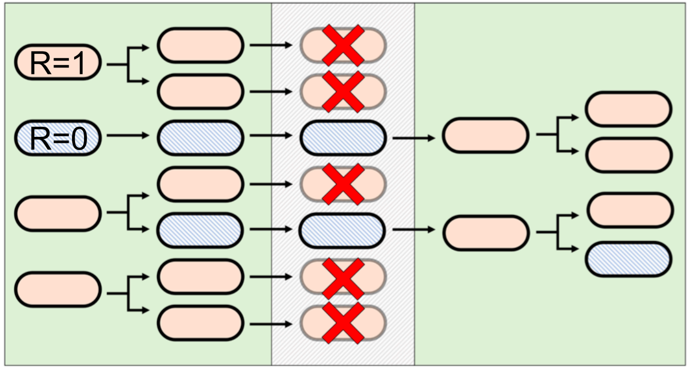

Question: optimal *u* given *p*?

Model assumptions:

- 2 states R: growing (R=1) / dormant (R=0)
- 2 environments E: antibiotic (E=+) / no antibiotic (E=-)
- survival/reproduction per generation f(R,E)

	E = +	E = -
R = 0	1	1
R = 1	0	2

- probability for antibiotic (E=+) : *p*
- probability to be dormant (R=0): *u*

(per generation)

Bacterial persistence: elementary model

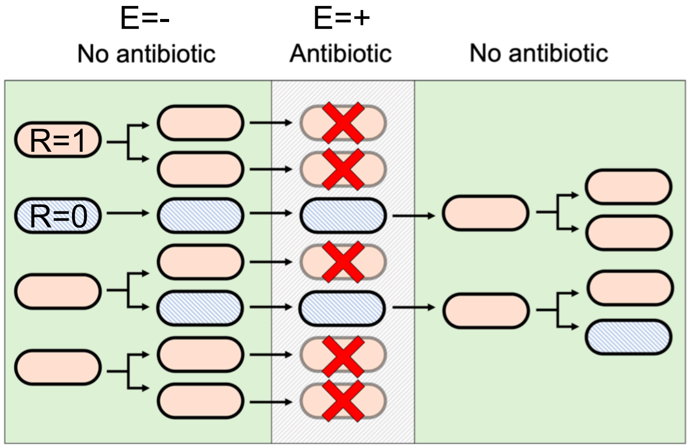

Question: optimal *u* given *p*?

Model assumptions:

- 2 states R: growing (R=1) / dormant (R=0)
- 2 environments E: antibiotic (E=+) / no antibiotic (E=-)
- survival/reproduction per generation f(R,E)

	E = +	E = -
R = 0	1	1
R = 1	0	2

- probability for antibiotic (E=+) : *p*
- probability to be dormant (R=0): *u*

(per generation)

Meta question: optimal in what sense?

Bacterial persistence: elementary model

Two convenient limits: (1) Large population
(2) Long time

	E = + p	E = - $1- p$
R = 0 u	1	1
R = 1 $1-u$	0	2

Bacterial persistence: elementary model

Two convenient limits: (1) Large population
(2) Long time

In one generation, given N_t cells at generation *t*, a fraction *u* is dormant (R=0) and *1-u* is growing (R=1):

if antibiotic (E=+): $N_{t+1} = A_+ N_t \qquad A_+ = u$

if no antibiotic (E=-): $N_{t+1} = A_- N_t \qquad A_- = u + 2(1-u) = 2 - u$

	E = + *p*	E = - *1- p*
R = 0 *u*	1	1
R = 1 *1-u*	0	2

Bacterial persistence: elementary model

Two convenient limits: (1) Large population
(2) Long time

In one generation, given N_t cells at generation *t*, a fraction *u* is dormant (R=0) and *1-u* is growing (R=1):

if antibiotic (E=+): $N_{t+1} = A_+ N_t$ $A_+ = u$

if no antibiotic (E=-): $N_{t+1} = A_- N_t$ $A_- = u + 2(1-u) = 2 - u$

Over *T* generations, a fraction *p* of generations with antibiotics (E=+) and *1-p* without (E=-):

$$N_T = (A_+)^{pT}(A_-)^{(1-p)T} N_0 = e^{\Lambda T} N_0$$

$$\Lambda = p \ln A_+ + (1-p)\ln A_- = p \ln u + (1-p)\ln(2-u)$$

	E = + *p*	E = - *1- p*
R = 0 *u*	1	1
R = 1 *1-u*	0	2

Bacterial persistence: elementary model

Two convenient limits: (1) Large population
(2) Long time

In one generation, given N_t cells at generation *t*, a fraction *u* is dormant (R=0) and *1-u* is growing (R=1):

if antibiotic (E=+): $N_{t+1} = A_+ N_t$ $\quad A_+ = u$

if no antibiotic (E=-): $N_{t+1} = A_- N_t$ $\quad A_- = u + 2(1-u) = 2-u$

Over *T* generations, a fraction *p* of generations with antibiotics (E=+) and *1-p* without (E=-):

$$N_T = (A_+)^{pT}(A_-)^{(1-p)T} N_0 = e^{\Lambda T} N_0$$

$$\Lambda = p \ln A_+ + (1-p)\ln A_- = p \ln u + (1-p)\ln(2-u)$$

	E = + *p*	E = - *1- p*
R = 0 *u*	1	1
R = 1 *1-u*	0	2

Optimal *u* $(\max \Lambda)$:

$$u = \begin{cases} 2p\,, & \text{if } 0 < p \le 1/2. \\ 1, & \text{if } 1/2 < p \le 1. \end{cases}$$

Bacterial persistence: elementary model

Two convenient limits: (1) Large population
(2) Long time

In one generation, given N_t cells at generation *t*, a fraction *u* is dormant (R=0) and *1-u* is growing (R=1):

if antibiotic (E=+): $N_{t+1} = A_+ N_t$ $\quad A_+ = u$

if no antibiotic (E=-): $N_{t+1} = A_- N_t$ $\quad A_- = u + 2(1-u) = 2 - u$

	E = + *p*	E = - *1- p*
R = 0 *u*	1	1
R = 1 *1-u*	0	2

Over *T* generations, a fraction *p* of generations with antibiotics (E=+) and *1-p* without (E=-):

$$N_T = (A_+)^{pT}(A_-)^{(1-p)T} N_0 = e^{\Lambda T} N_0$$

$$\Lambda = p \ln A_+ + (1-p)\ln A_- = p \ln u + (1-p)\ln(2-u)$$

Optimal *u* $(\max \Lambda)$**:**

$$u = \begin{cases} 2p\ , & \text{if } 0 < p \leq 1/2. \\ 1, & \text{if } 1/2 < p \leq 1. \end{cases}$$

Conclusion: The optimal fraction of persisters *u* depends on the environmental uncertainty *p*

Generalization, including sensing

Phenotypic switching: random transitions between phenotypes independent of the environment

Sensing: switch to a new phenotype R depending on a cue S correlated to the environment E

Generalization, including sensing

Phenotypic switching: random transitions between phenotypes independent of the environment

Sensing: switch to a new phenotype R depending on a cue S correlated to the environment E

Model assumptions:

- n states R
- M environmental states E, probability $p(E)$
- multiplication factor $f(R,E)$
- switching probability $u(R|S)$ where S is a cue
- probability $q(S|E)$ for S given E

Generalization, including sensing

Phenotypic switching: random transitions between phenotypes independent of the environment

Sensing: switch to a new phenotype R depending on a cue S correlated to the environment E

Model assumptions:

- n states R
- M environmental states E, probability p(E)
- multiplication factor f(R,E)
- switching probability u(R|S) where S is a cue
- probability q(S|E) for S given E

Dynamics:

- 1 generation: $N_{t+1} = A(E,S)N_t$

$$A(E,S) = \sum_R f(R,E)u(R|S)$$

- T generations: $N_T = e^{\Lambda T} N_0$

$$\Lambda = \sum_{S,E} q(S|E)p(E) \ln A(E,S)$$

Generalization, including sensing

Phenotypic switching: random transitions between phenotypes independent of the environment

Sensing: switch to a new phenotype R depending on a cue S correlated to the environment E

Model assumptions:

- n states R
- M environmental states E, probability p(E)
- multiplication factor f(R,E)
- switching probability u(R|S) where S is a cue
- probability q(S|E) for S given E

Dynamics:

- 1 generation: $N_{t+1} = A(E,S)N_t$

$$A(E,S) = \sum_R f(R,E)u(R|S) \quad = \langle f(R,E)\rangle_R$$

$$\sum_R u(R|S) = 1$$

- T generations: $N_T = e^{\Lambda T} N_0$

$$\Lambda = \sum_{S,E} q(S|E)p(E)\ln A(E,S) \quad = \langle \ln A(E,S)\rangle_{E,S}$$

$$\sum_{S,E} q(S|E)p(E) = 1$$

"Fitness" = long-term growth rate $\Lambda = \langle \ln(\langle f(R,E)\rangle_R)\rangle_{E,S}$

average over phenotypes within a generation (arithmetic mean) → $\langle \cdot \rangle_R$

average over environments across generations (geometric mean) → $\langle \cdot \rangle_{E,S}$

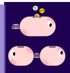

Geometric *vs* arithmetic means

Growth is a multiplicative process

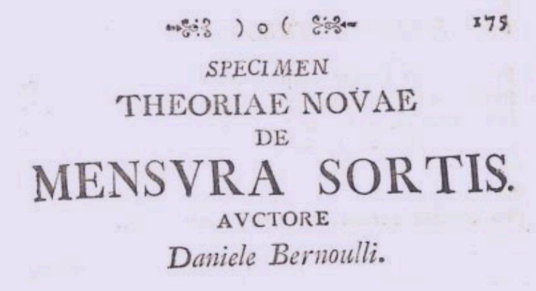

175

SPECIMEN
THEORIAE NOVAE
DE
MENSVRA SORTIS.
AVCTORE
Daniele Bernoulli.

D Bernoulli, *Exposition of a new theory on the measurement of risk* (1738)

Geometric *vs* arithmetic means

Growth is a multiplicative process

Didactic example:

A=2 with *p*=1/2 or *A*=1/3 with *p*=1/2

arithmetic mean = $(2+1/3)/2 > 1$

geometric mean = $2^{1/2}(1/3)^{1/2} < 1$

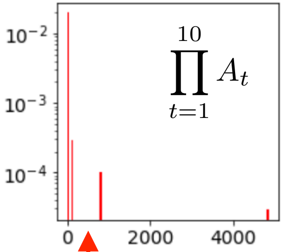

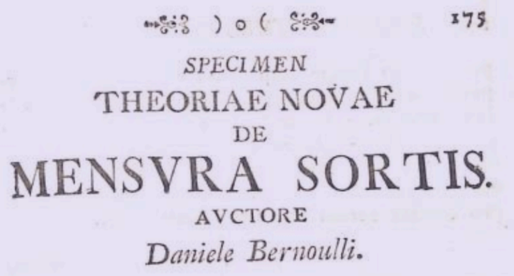
175

SPECIMEN
THEORIAE NOVAE
DE
MENSVRA SORTIS.
AVCTORE
Daniele Bernoulli.

D Bernoulli, *Exposition of a new theory on the measurement of risk* (1738)

Geometric *vs* arithmetic means

Growth is a multiplicative process

Didactic example:

A=2 with *p*=1/2 or *A*=1/3 with *p*=1/2

arithmetic mean = $(2+1/3)/2 > 1$

geometric mean = $2^{1/2}(1/3)^{1/2} < 1$

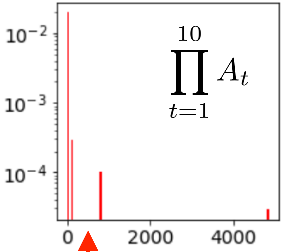

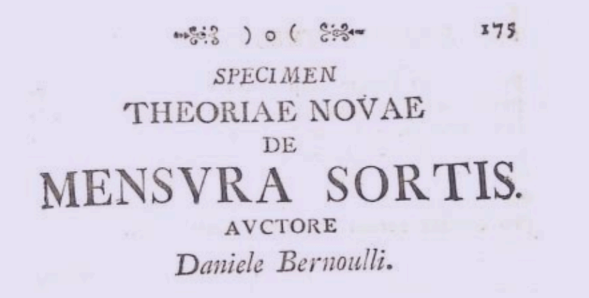

D Bernoulli, *Exposition of a new theory on the measurement of risk* (1738)

Back to the simple model of bacterial persistence:

	E = + *p*	E = - *1- p*
R = 0 *u*	1	1
R = 1 *1-u*	0	2

Growth per generation:

$$A(E=+) = u$$

$$A(E=-) = u + 2(1-u)$$

Geometric *vs* arithmetic means

Growth is a multiplicative process

Didactic example:

A=2 with *p*=1/2 or *A*=1/3 with *p*=1/2

arithmetic mean = $(2+1/3)/2 > 1$

geometric mean = $2^{1/2}(1/3)^{1/2} < 1$

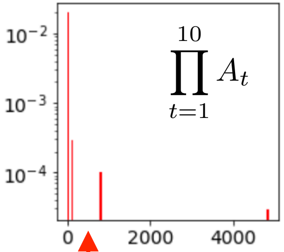

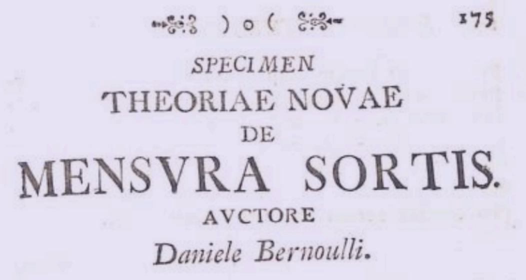
175

SPECIMEN
THEORIAE NOVAE
DE
MENSVRA SORTIS.
AVCTORE
Daniele Bernoulli.

D Bernoulli, *Exposition of a new theory on the measurement of risk* (1738)

Back to the simple model of bacterial persistence:

	E = + *p*	E = - *1- p*
R = 0 *u*	1	1
R = 1 *1-u*	0	2

Growth per generation:

$$A(E=+) = u$$

$$A(E=-) = u + 2(1-u)$$

Optimal arithmetic mean:

$$\max\langle A(E)\rangle_E \qquad u = \begin{cases} 0, & \text{if } 0 < p \le 1/2 \\ 1, & \text{if } 1/2 < p \le 1 \end{cases}$$

***u*=0 is very risky:** extinction when E=+ occurs

Geometric *vs* arithmetic means

Growth is a multiplicative process

Didactic example:

A=2 with *p*=1/2 or *A*=1/3 with *p*=1/2

arithmetic mean = $(2+1/3)/2 > 1$

geometric mean = $2^{1/2}(1/3)^{1/2} < 1$

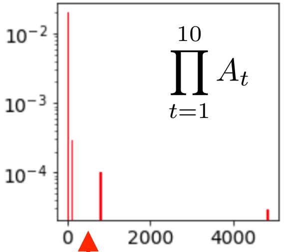

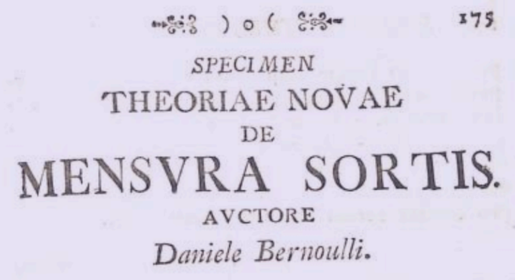
) o (175

SPECIMEN
THEORIAE NOVAE
DE
MENSVRA SORTIS.
AVCTORE
Daniele Bernoulli.

D Bernoulli, *Exposition of a new theory on the measurement of risk* (1738)

Back to the simple model of bacterial persistence:

	E = + *p*	E = - *1- p*
R = 0 *u*	1	1
R = 1 *1-u*	0	2

Growth per generation:

$$A(E=+) = u$$

$$A(E=-) = u + 2(1-u)$$

Optimal arithmetic mean:

$$\max\langle A(E)\rangle_E \qquad u = \begin{cases} 0, & \text{if } 0 < p \leq 1/2 \\ 1, & \text{if } 1/2 < p \leq 1 \end{cases}$$

***u*=0 is very risky:** extinction when E=+ occurs

Optimal geometric mean:

$$\max\langle \ln A(E)\rangle_E \qquad u = \begin{cases} 2p, & \text{if } 0 < p \leq 1/2 \\ 1, & \text{if } 1/2 < p \leq 1 \end{cases}$$

Value and cost of sensing information

Long-term growth rate:

$$\Lambda = \langle \ln A(S,E) \rangle_{S,E} = \sum_{S,E} q(S|E) p(E) \ln A(S,E)$$

$$A(S,E) = \langle f(R,E) \rangle_R = \sum_R f(R,E) u(R|S)$$

Value and cost of sensing information

Long-term growth rate: $\Lambda = \langle \ln A(S,E) \rangle_{S,E} = \sum_{S,E} q(S|E)p(E) \ln A(S,E)$ $\quad A(S,E) = \langle f(R,E) \rangle_R = \sum_R f(R,E)u(R|S)$

Simplifying assumption: only one phenotype $R(E)$ can survive for each environmental state E (Kelly model)

$f(R,E) = f(E)$ if $R = R(E)$ and 0 otherwise $\implies A(E,S) = f(E)u(E|S)$

Value and cost of sensing information

Long-term growth rate: $\Lambda = \langle \ln A(S,E) \rangle_{S,E} = \sum_{S,E} q(S|E) p(E) \ln A(S,E)$ $\quad A(S,E) = \langle f(R,E) \rangle_R = \sum_R f(R,E) u(R|S)$

Simplifying assumption: only one phenotype $R(E)$ can survive for each environmental state E (Kelly model)

$f(R,E) = f(E)$ if $R = R(E)$ and 0 otherwise $\implies A(E,S) = f(E) u(E|S)$

Optimum without sensing:

$$\Lambda_0^* = \sum_E p(E) \ln f(E) - H(p) \qquad H(p) = -\sum_E p(E) \ln p(E)$$

growth limited by the entropy of the environment

Value and cost of sensing information

Long-term growth rate: $\Lambda = \langle \ln A(S,E) \rangle_{S,E} = \sum_{S,E} q(S|E)p(E) \ln A(S,E)$ $\quad A(S,E) = \langle f(R,E) \rangle_R = \sum_R f(R,E)u(R|S)$

Simplifying assumption: only one phenotype $R(E)$ can survive for each environmental state E (Kelly model)

$f(R,E) = f(E)$ if $R = R(E)$ and 0 otherwise $\implies A(E,S) = f(E)u(E|S)$

Optimum without sensing:

$$\Lambda_0^* = \sum_E p(E) \ln f(E) - H(p) \qquad H(p) = -\sum_E p(E) \ln p(E)$$

growth limited by the entropy of the environment

Optimum with sensing, i.e. given q(S|E)

$$\Lambda_q^* - \Lambda_0^* = \sum_{S,E} q(S|E)p(E) \ln q(S|E) = I(S,E)$$

growth increase given by the mutual information between the signal and the environment

Value and cost of sensing information

Long-term growth rate: $\Lambda = \langle \ln A(S,E) \rangle_{S,E} = \sum_{S,E} q(S|E) p(E) \ln A(S,E)$ $\quad A(S,E) = \langle f(R,E) \rangle_R = \sum_R f(R,E) u(R|S)$

Simplifying assumption: only one phenotype $R(E)$ can survive for each environmental state E (Kelly model)

$f(R,E) = f(E)$ if $R = R(E)$ and 0 otherwise $\implies A(E,S) = f(E) u(E|S)$

Optimum without sensing:

$$\Lambda_0^* = \sum_E p(E) \ln f(E) - H(p) \qquad H(p) = -\sum_E p(E) \ln p(E)$$

growth limited by the entropy of the environment

Optimum with sensing, i.e. given q(S|E)

$$\Lambda_q^* - \Lambda_0^* = \sum_{S,E} q(S|E) p(E) \ln q(S|E) = I(S,E)$$

growth increase given by the mutual information between the signal and the environment

Trade-off between value of information *I(S,E)* and cost *c* of the sensor: may be optimal to have no sensor

Value and cost of sensing information

Long-term growth rate: $\Lambda = \langle \ln A(S,E) \rangle_{S,E} = \sum_{S,E} q(S|E)p(E) \ln A(S,E)$ $\quad A(S,E) = \langle f(R,E) \rangle_R = \sum_R f(R,E)u(R|S)$

Simplifying assumption: only one phenotype $R(E)$ can survive for each environmental state E (Kelly model)

$f(R,E) = f(E)$ if $R = R(E)$ and 0 otherwise $\implies A(E,S) = f(E)u(E|S)$

Optimum without sensing:

$$\Lambda_0^* = \sum_E p(E) \ln f(E) - H(p) \qquad H(p) = -\sum_E p(E) \ln p(E)$$

growth limited by the entropy of the environment

Optimum with sensing, i.e. given q(S|E)

$$\Lambda_q^* - \Lambda_0^* = \sum_{S,E} q(S|E)p(E) \ln q(S|E) = I(S,E)$$

growth increase given by the mutual information between the signal and the environment

Trade-off between value of information *I(S,E)* and cost *c* of the sensor: may be optimal to have no sensor

Generalization & links with information theory: Covers & Thomas, *Information Theory*

Analogies with finance & beyond

Biology (**population**)	**Finance** (**capital**)
Individual	Currency unit
Environment $p(E)$	Market state
–	Investor
Phenotype decisions $u(R)$	Investment strategy
Multiplicative rate $f(R, E)$	Immediate return
Environmental cue $P(S\|E)$	Side information

Analogies with finance & beyond

Biology (population)	**Finance (capital)**
Individual	Currency unit
Environment $p(E)$	Market state
–	Investor
Phenotype decisions $u(R)$	Investment strategy
Multiplicative rate $f(R, E)$	Immediate return
Environmental cue $P(S\|E)$	Side information

→

Major difference

Biology: each cell processes information

Finance: one investor centralizes information

Analogies with finance & beyond

Biology (population)	Finance (capital)
Individual	Currency unit
Environment $p(E)$	Market state
–	Investor
Phenotype decisions $u(R)$	Investment strategy
Multiplicative rate $f(R, E)$	Immediate return
Environmental cue $P(S\|E)$	Side information

Major difference

Biology: each cell processes information

Finance: one investor centralizes information

Implication 1: if one sensor per cell, cell-to-cell heterogeneity in perceived signals

Different averages: $\Lambda = \langle \ln \langle f(E,R) \rangle_R \rangle_{S,E}$ < $\Lambda = \langle \ln \langle f(E,R) \rangle_{R,S} \rangle_E$

if S is centralized — *if S is distributed*

Analogies with finance & beyond

Biology (**population**)	**Finance** (**capital**)
Individual	Currency unit
Environment $p(E)$	Market state
–	Investor
Phenotype decisions $u(R)$	Investment strategy
Multiplicative rate $f(R,E)$	Immediate return
Environmental cue $P(S\|E)$	Side information

→

Major difference

Biology: each cell processes information

Finance: one investor centralizes information

Implication 1: if one sensor per cell, cell-to-cell heterogeneity in perceived signals

Different averages: $\Lambda = \langle \ln \langle f(E,R) \rangle_R \rangle_{S,E}$ < $\Lambda = \langle \ln \langle f(E,R) \rangle_{R,S} \rangle_E$

if S is centralized *if S is distributed*

Implication 2: what is optimal for a population may not be evolutionary stable

Possible conflict between levels of selection (tragedy of the commons)

No conflict in the models presented here but, more generally, optimal $\neq$ evolvable

Phenotypic vs genotypic information

'Standard model' of biological information processing

- survival/ reproduction (ξ) depends on the phenotype R_t
- the phenotype R_t depends on an inherited genotype γ_t
- the inherited genotype γ_t determines the transmitted genotype γ_{t+1}

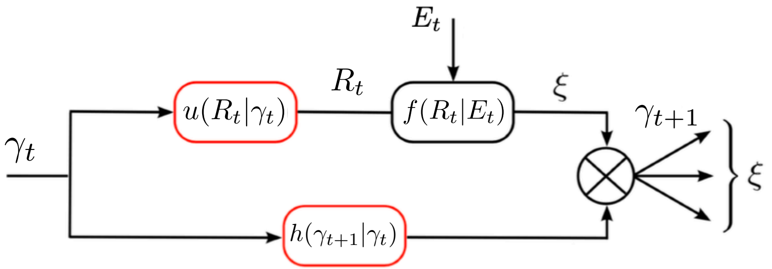

Phenotypic vs genotypic information

'Standard model' of biological information processing

- survival/ reproduction (ξ) depends on the phenotype R_t
- the phenotype R_t depends on an inherited genotype γ_t
- the inherited genotype γ_t determines the transmitted genotype γ_{t+1}

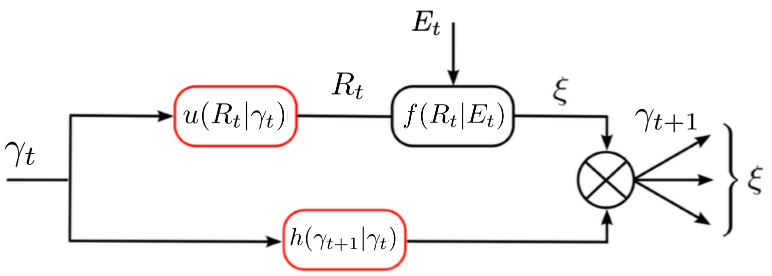

Two sources of stochasticity

- development $u(R_t|\gamma_t)$ (generalizing phenotypic switching)
- mutations $h(\gamma_{t+1}|\gamma_t)$

Phenotypic vs genotypic information

'Standard model' of biological information processing

- survival/ reproduction (ξ) depends on the phenotype R_t
- the phenotype R_t depends on an inherited genotype γ_t
- the inherited genotype γ_t determines the transmitted genotype γ_{t+1}

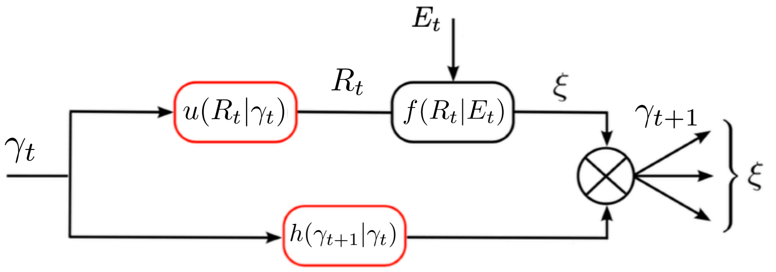

Two sources of stochasticity

- development $u(R_t|\gamma_t)$ (generalizing phenotypic switching)
- mutations $h(\gamma_{t+1}|\gamma_t)$

Nature of optimal stochasticity in uncertain environments?

Depends on:
- the variance of environmental fluctuations
- the correlation between successive environments

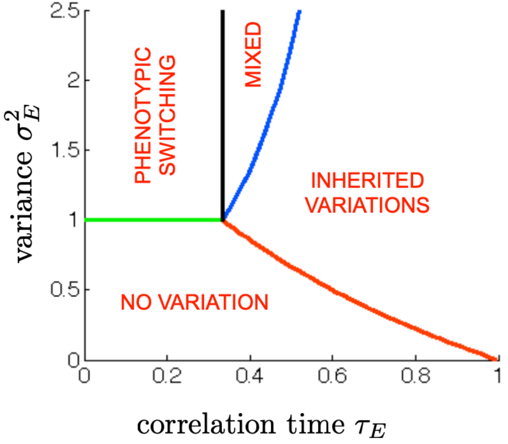

Phenotypic vs genotypic information

'Standard model' of biological information processing

- survival/ reproduction (ξ) depends on the phenotype R_t
- the phenotype R_t depends on an inherited genotype γ_t
- the inherited genotype γ_t determines the transmitted genotype γ_{t+1}

Two sources of stochasticity

- development $u(R_t|\gamma_t)$ (generalizing phenotypic switching)
- mutations $h(\gamma_{t+1}|\gamma_t)$

Nature of optimal stochasticity in uncertain environments?

Depends on:
- the variance of environmental fluctuations
- the correlation between successive environments

Extension: other mechanisms to generate and transmit variations

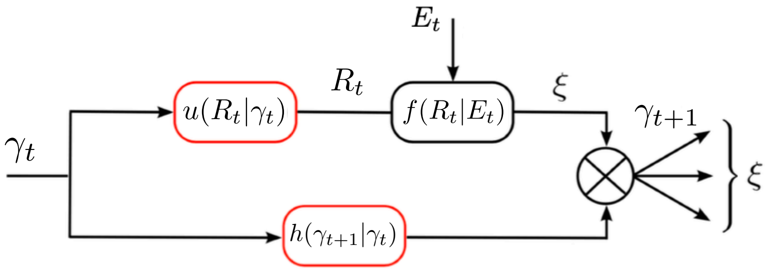

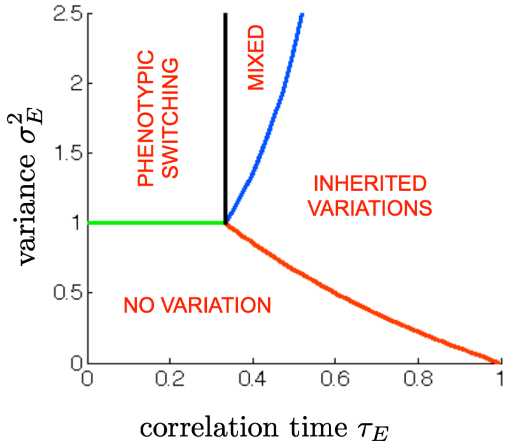

Summary and perspectives

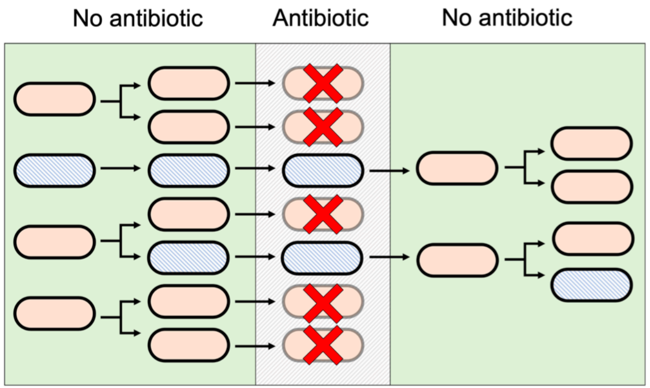

Example: Bacterial persistence

Adaptation to uncertain environments

- Long-term population-level adaptation
- Phenotypic switching versus sensing
- Phenotypic switching versus mutations

Summary and perspectives

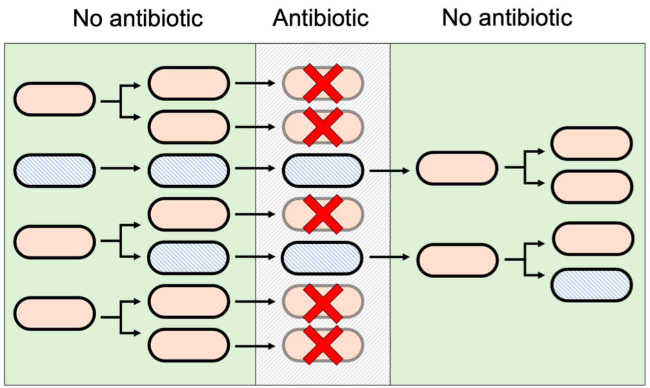

Example: Bacterial persistence

Adaptation to uncertain environments

- Long-term population-level adaptation
- Phenotypic switching versus sensing
- Phenotypic switching versus mutations

Analogy with games and finance

- Bet-hedging / portfolio diversification
- Key difference: level of information processing

Summary and perspectives

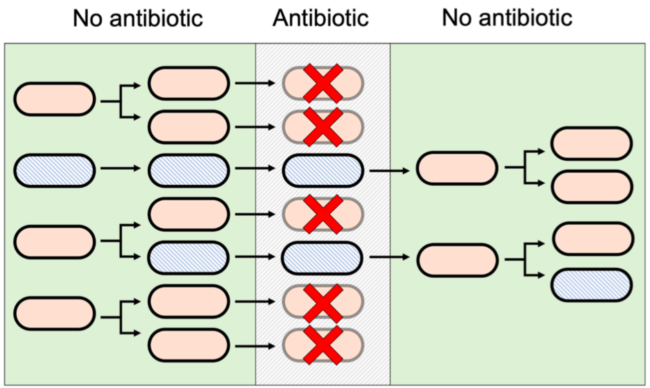

Example: Bacterial persistence

Adaptation to uncertain environments

- Long-term population-level adaptation
- Phenotypic switching versus sensing
- Phenotypic switching versus mutations

Analogy with games and finance

- Bet-hedging / portfolio diversification
- Key difference: level of information processing

Mathematical formalism

- Geometric versus arithmetic means
- Quantifying information with entropies

Summary and perspectives

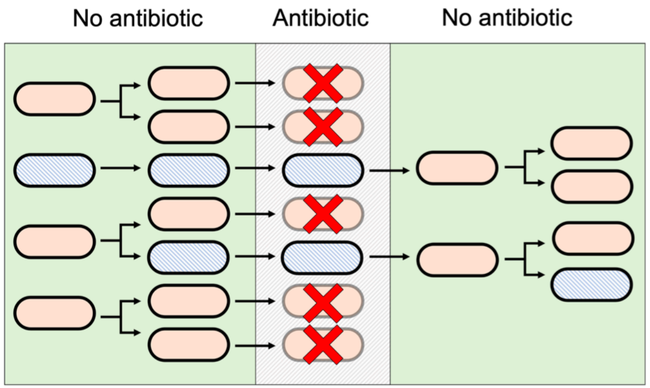

Example: Bacterial persistence

Adaptation to uncertain environments

- Long-term population-level adaptation
- Phenotypic switching versus sensing
- Phenotypic switching versus mutations

Analogy with games and finance

- Bet-hedging / portfolio diversification
- Key difference: level of information processing

Mathematical formalism

- Geometric versus arithmetic means
- Quantifying information with entropies

Model based on several assumptions:

- Long-term growth rate (many generations)
- Large population (no extinction)
- Environment independent of population dynamics

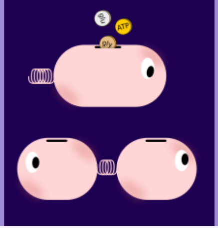

Cells in the face of uncertainty
part II

D. Lacoste

INRAE Département MathNum EURIP LEARNING PLANET INSTITUTE

Outline of the talk

1. Tradeoff in optimal gambling strategies

2. Adaptive strategies in gambling

3. Tradeoff for phenotypic switching of populations in varying environments

Kelly’s formula in popular culture

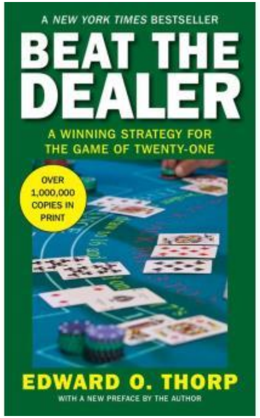

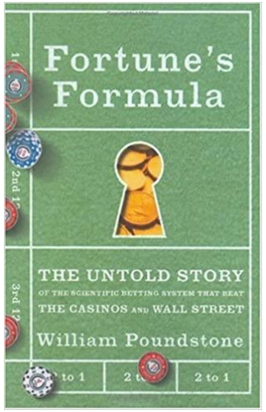

From card counting method in blackjack. .. .. to investments on the stock market

A new interpretation of information rate, Kelly J. L. J. (1956)

Kelly’s model as a resource allocation problem

Gambler **Bookmaker**

Bets

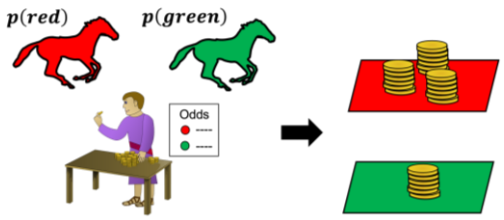

Odds

Constraints : $\sum_{x=1}^{M} b_x = 1$ and $r_x := \frac{1}{o_x}$ with $\sum_{x=1}^{M} r_x = 1$ for fair odds

Dynamics : winning horse x is chosen with probability p_x

Then capital is updated : $C_{t+1} = \frac{\mathrm{b}_x}{\mathrm{r}_x} C_t$

Long term growth rate

Log-Capital $\log\text{-cap}(t) = \sum_{\tau=1}^{t} \log\left(\frac{\mathrm{b}_{x_\tau}}{\mathrm{r}_{x_\tau}}\right)$

by the law of large numbers : $\frac{\log\text{-cap}(t)}{t} \xrightarrow[t\to\infty]{} \mathbb{E}\left[\log\left(\frac{\mathrm{b}_x}{\mathrm{r}_x}\right)\right]$

Optimization of the long term growth rate (Kelly's optimal strategy)

$$\langle W \rangle = \mathbb{E}\left[\log\left(\frac{\mathrm{b}_x}{\mathrm{r}_x}\right)\right] = D_{KL}\left(\mathbf{p}||\mathbf{r}\right) - D_{KL}\left(\mathbf{p}||\mathbf{b}\right)$$

This is maximum when $b_x = p_x$ and at this point $\langle W^* \rangle = D_{KL}\left(\mathbf{p}||\mathbf{r}\right) \geq 0$

The gambler makes money when he/she has better knowledge of the winning probabilities than the bookie

Evolution of the capital of the gambler

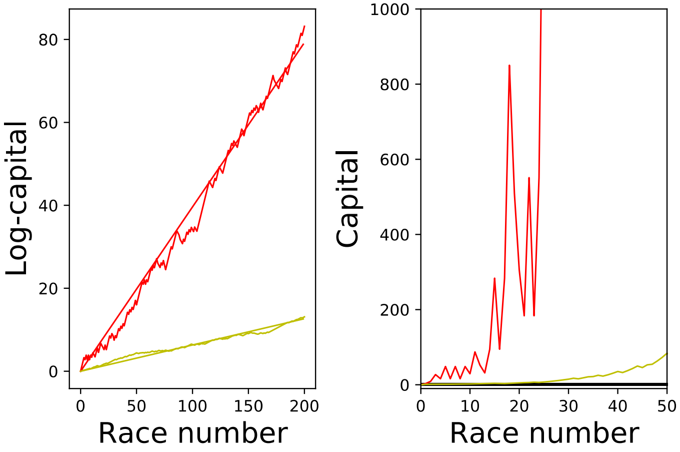

- Kelly’s strategy dominates on long times all non-optimal strategies

- A general trade-off between the maximization of the growth rate and the minimization of risky fluctuations ?

L. Dinis, J. Unterberger, D. L., Eur. Phys. Lett. (2020)

How to define risk ?

By the central limit theorem :

$$\frac{1}{\sigma_W \sqrt{t}} \left(\log \frac{C_t}{C_0} - t\langle W \rangle \right) \xrightarrow[t\to\infty]{} \mathcal{N}(0,1) \text{ normal law}$$

where $\sigma_W^2 = \mathrm{Var}\left[\log \left(\frac{\mathrm{b}_x}{\mathrm{r}_x} \right) \right]$ is the volatility

The volatility is not the best measure of risk but it leads to tractable calculations

In practice, risk is relevant at intermediate time scales $t \ll (\sigma_W / \langle W \rangle)^2$

Risk free strategy

Note that the strategy $b_x = r_x$ has $\sigma_W = 0$ and $\langle W \rangle = 0$

Objective function :

$$J = \alpha \langle W \rangle - (1-\alpha)\sigma_W + \lambda \sum_x b_x$$

- Interpolates between maximization of growth rate for α=1 and the minimization of the fluctuations when α=0
- The optimal solution is parametrized by α, which is a risk aversion parameter.
- Similarities with Markowitz portfolio theory

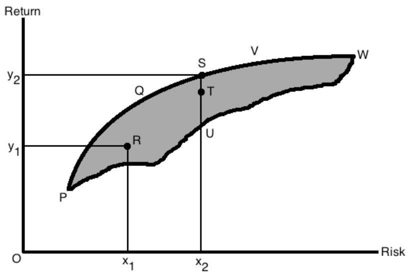

Markowitz H. (1952)

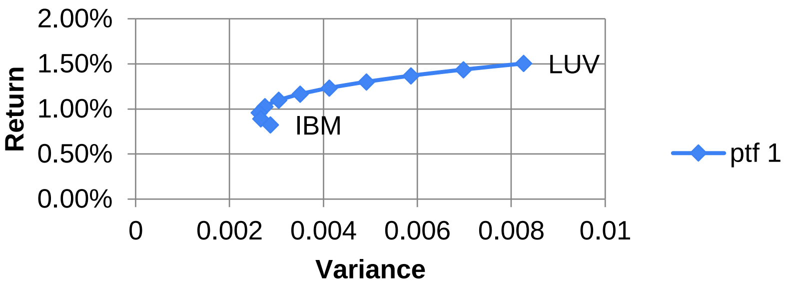

Data from Wharton School of Finance

The efficient border for two horses problem

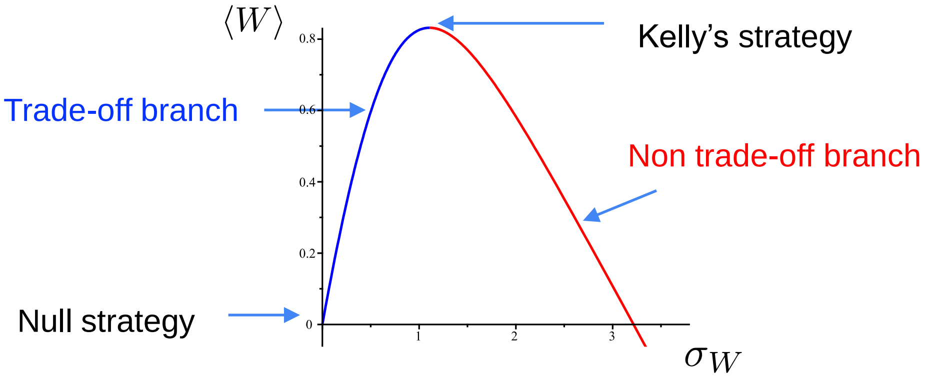

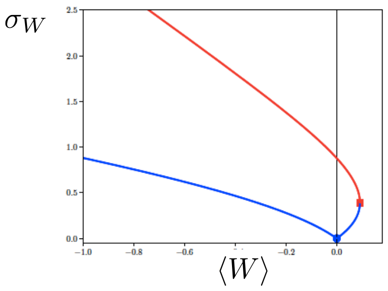

For p<r :

In the $\langle W \rangle \geq 0$ region, $\dfrac{d\sigma_W}{d\langle W \rangle} = \dfrac{\sigma}{p-b}$ becomes infinite near Kelly's strategy

but non-zero near the null strategy where :

$$\frac{d\sigma_W}{d\langle W \rangle} = \frac{1}{\gamma_c} = \frac{\sigma}{|p-r|} \quad \text{and} \quad \frac{d^2\sigma_W}{d\langle W \rangle^2} = \frac{r(1-r)}{\sigma^2\gamma_c^3} > 0$$

Beyond 2 horses : numerical optimization

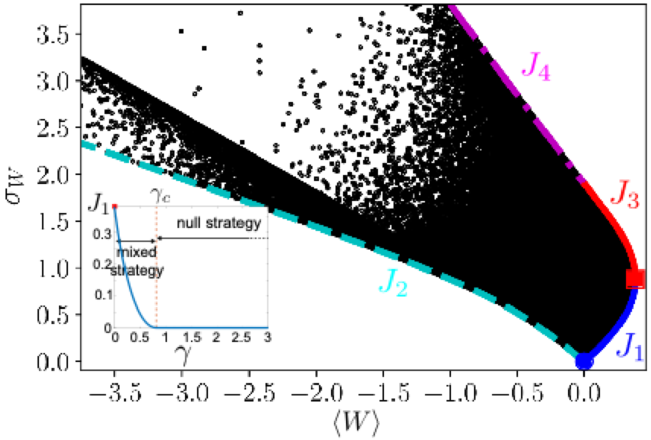

In practice, the numerical optimization of the objective function
relies on simulated annealing or Karush-Kuhn-Tücker (KKT) algorithms.

Game theoretic formulation

- Worst possible case for the gambler corresponds to minimization of

$$\Psi(\mathbf{p}) = \langle W(\mathbf{p}, \mathbf{b}^{\text{KELLY}}) \rangle - \lambda \sum_x \mathrm{p}_x$$

$$\mathrm{p}_x = \mathrm{p}_x^* = \frac{\mathrm{r}_x}{\sum_x \mathrm{r}_x}$$

- The general growth rate is

$$\langle W(\mathbf{p}, \mathbf{b}) \rangle = D_{KL}(\mathbf{p}||\mathbf{p}^*) - D_{KL}(\mathbf{p}||\mathbf{b}) + V$$

R. Pugatch et al., (2014)

$D_{KL}(\mathbf{p}||\mathbf{p}^*)$ *pessimistic surprise* : things are not as bad as one would think

$-D_{KL}(\mathbf{p}||\mathbf{b})$ *gambler's regret* : gambler plays sub-optimally

V *value of the game : V<0 for unfair odds, V>0 for super-fair odds*

Non-diagonal odds

- Now, the growth rate is : $$\langle W(\mathbf{p},\mathbf{b})\rangle = \sum_x \mathrm{p}_x \ln\left(\sum_y \mathrm{o}_{xy}\mathrm{b}_y\right)$$

- When the odds matrix is invertible $\mathrm{r} = \mathrm{o}^{-1}$ and simplex preserving (fully mixing game)

Optimal bets : $\mathrm{b}^*_x = \sum_y \Omega_{xy}\mathrm{p}_y$ with $\Omega_{xy} = \frac{\mathrm{r}_{xy}}{\sum_l \mathrm{r}_{ly}}$

Optimal environment : $\mathrm{p}^*_x = \frac{\sum_l \mathrm{r}_{lx}}{\sum_{xy}\mathrm{r}_{xy}}$

$(\mathrm{b}^*_x, \mathrm{p}^*_x)$ defines a Nash equilibrium

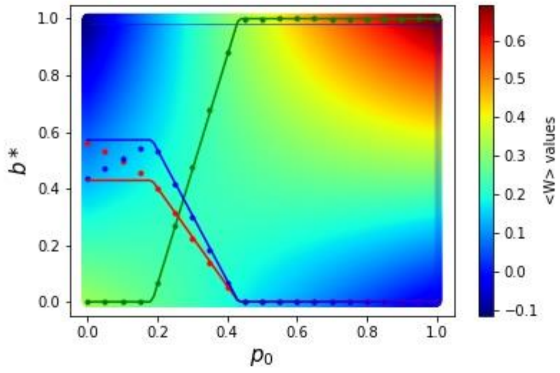

S. Cavallero, (2023)

Mean-variance trade-offs

- For fair odds, assuming $\langle W \rangle \geq 0$ with q the pdf such that $q_x := r_x / p_x$

$$\sigma_W \geq \frac{\langle W \rangle}{\sigma_q}$$

L. Dinis et al., EPL (2020)

- For non-fair odds with $\langle q \rangle = \sum_x r_x \neq 1$ and $V = -\log \sum_x r_x$

$$\sigma_W \geq \frac{|V - \langle W \rangle|}{\sigma_q} \langle q \rangle$$

General trade-off between growth rate and risk

Similar tradeoff in the thermodynamics of non-equilibrium systems A. Barato et al., (2015)

Numerical illustration

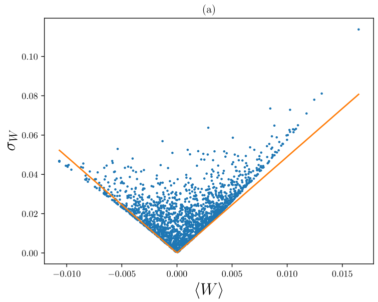

Non-diagonal fair odds

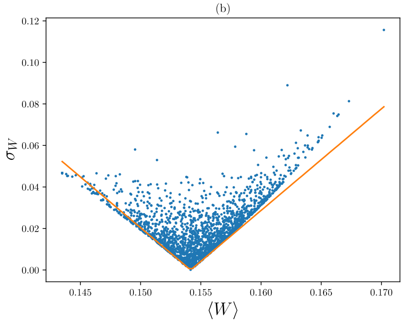

Non-diagonal super-fair odds

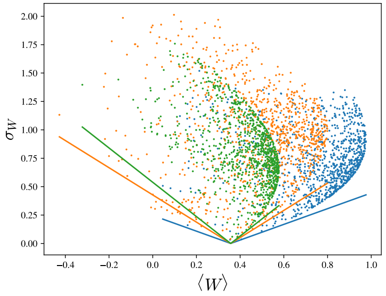

2. Adaptive strategies in gambling

- So far, we assumed the gambler knows the probabilities of winning horses,

In practice the gambler does not know this, *he/she must learn it* !

- This learning can be modeled using *Laplace's rule of succession* (equivalent to Bayesian inference)

$$b_x^{t+1} = \frac{n_x^t + 1}{t + M}$$

E. T. Jaynes, 2003

for t uncorrelated races and M horses

- Gambler's regret : the difference between the actual growth rate and the one of the optimal strategy :

$$\Delta(t) = \text{log-cap}^{\text{Kelly}}(t) - \text{log-cap}(t)$$

The learning time and the gambler's regret

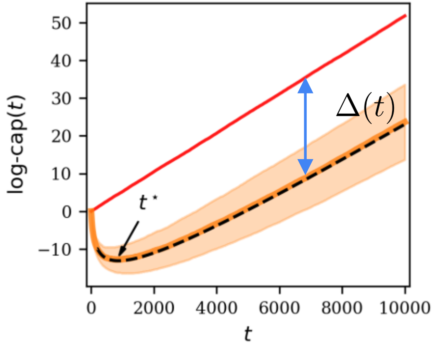

A. Despons et al. (2022)

Asymptotic regret : $\langle \Delta \rangle (t) = \langle \Delta \rangle (t_0) + \frac{M-1}{2} \log \frac{t}{t_0 + 1}$

Learning time : $t^\star = \frac{M-1}{2} \frac{1}{D_{KL}(\mathbf{p} \| \mathbf{r})}$

represents a limit on the characteristic time of variation of the environment

3. Trade-off for phenotypic switching of populations in varying environments

Bet-hedging and dormancy

Eriophyllum lanosum, plant from western USA desert

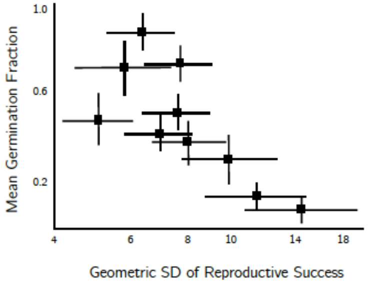

D. Venable (2007)

Fraction of seeds which germinated vs. standard deviation of reproductive success

Diversification (bet-hedging) is a universal adaptation strategy to an uncertain environnement

Seed bank: some seeds stay dormant to protect from harsch environments

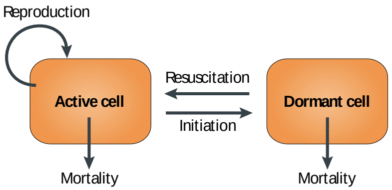

J. Lennon (2011)

Ecology

Biodiversity as insurance: from concept to measurement and application

Michel Loreau[1*], Matthieu Barbier[1], Elise Filotas[2], Dominique Gravel[3], Forest Isbell[4], Steve J. Miller[5], Jose M. Montoya[1], Shaopeng Wang[6], Raphaël Aussenac[7], Rachel Germain[8], Patrick L. Thompson[8], Andrew Gonzalez[9] and Laura E. Dee[10]

Microbial seed banks: the ecological and evolutionary implications of dormancy

Jay T. Lennon‡ and Stuart E. Jones*§*

Gambling/finance	Biology/ecology
Currency unit	Individual
Race result/market state	Environment
Bets/investment	Phenotype switching
Races	Environmental events
Odds	Reproduction rate
Capital growth rate	Population growth rate
Probability of bankruptcy	Extinction probability

- Sub-populations of two phenotypes growing in two environments $\frac{d}{dt}\mathbf{N}(t) = M_{S_i}\mathbf{N}(t)$ for $i \in \{1,2\}$

$$M_{S_1} = \begin{pmatrix} k_{A1} - \pi_1 & \pi_2 \\ \pi_1 & k_{B1} - \pi_2 \end{pmatrix} \text{ and } M_{S_2} = \begin{pmatrix} -\pi_1 + k_{A2} & \pi_2 \\ \pi_1 & k_{B2} - \pi_2 \end{pmatrix}.$$

- Gambling problem was *scalar*, this one is *vectorial.* Explicit results only in some limits

Ex: adiabatic limit E. Kussel, S. Leibler (2005)

Optimal condition is $\pi_i = k_i$ The analog of Kelly's strategy

- So far, we focused on long term growth rate (*infinite horizon*) but populations are finite and may go extinct in a finite time (*finite horizon*)

$$\mathrm{Var}(\Lambda) = \lim_{t\to\infty} t\mathrm{Var}(\Lambda_t)$$ is the equivalent of the volatility

Trade-off between growth and extinction probability

Fluctuations of growth rate

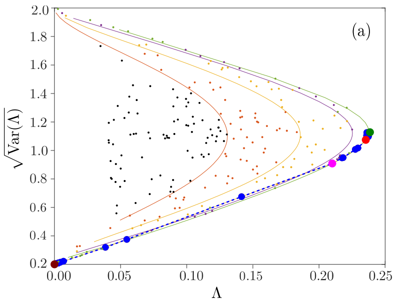

average growth rate

L. Dinis et al. (2022)

In the region of fast growth, it is advantageous for a population to trade growth for less risky fluctuations

Risk may be measured by growth rate fluctuations or extinction probability

Growth inhibition by antibiotics

- Most antibiotics do not kill cells directly but rather inhibit molecules involved in key cellular processes
- Risk may be measured by the fraction of inhibited molecules
- Risk correlates with pre-exposure growth rate and increases with the exposure to the drug

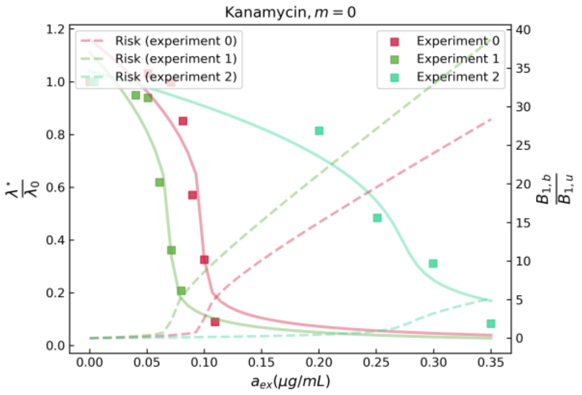

Economic principles of cell biology

i. When facing uncertainty, bet-hedging is a generic adaptation strategy for cells

Simplest form of this strategy is Kelly’s gambling

ii. There is a general trade-off between growth rate and risk exposure

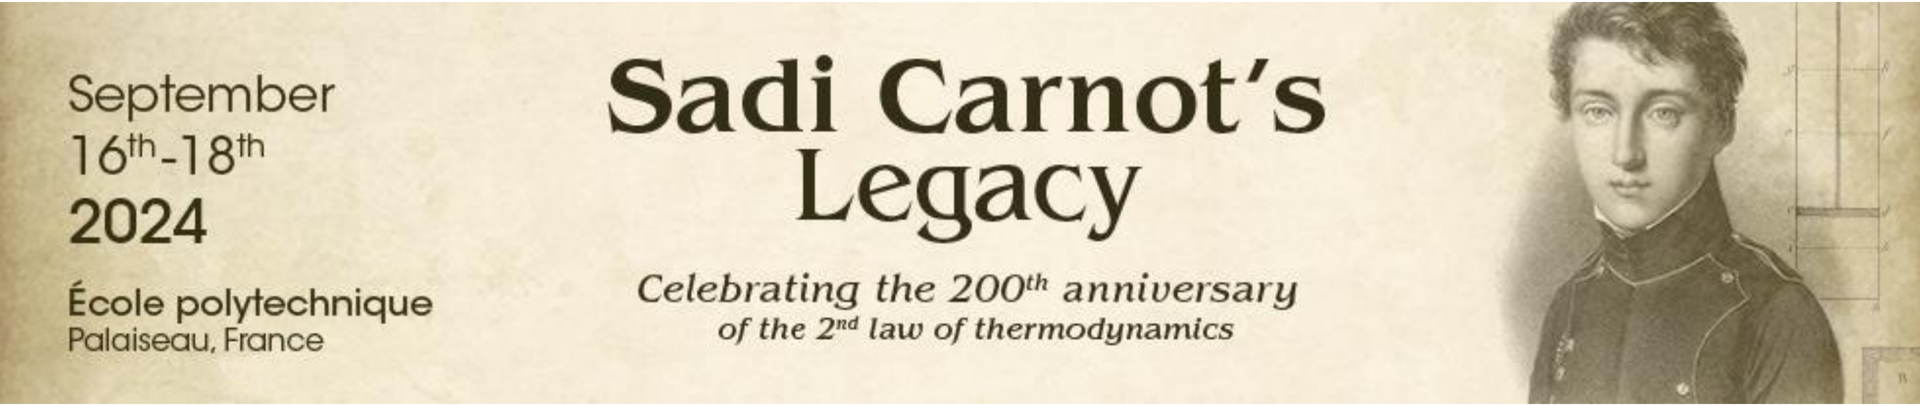

« Sur la puissance motrice du feu et sur les machines propres à développer cette puissance » (**1824**)

Acknowledgements

L. Dinis,
Universitad Complutense
Madrid

J. Unterberger,
Université de Lorraine

L. Peliti
Université de Naples

A. Despons

B. Ledoux