INFO-F-409
Learning dynamics

Evolutionary game theory, stochastic dynamics and the origins of co-operation

T. Lenaerts and A. Nowé
MLG, Université Libre de Bruxelles and
AI-lab, Vrije Universiteit Brussel

1

Summary

- What? Why?
- Rational choice
- Strategic games
- Nash Equilibrium
- Best
- Dominance
- Mixed strategies
- Mixed-strategy Nash Equilibria
- Support finding
- Lemke-Howson algorithm
- Extensive-form games
- sub-game perfect equilibrium
- Simultaneous moves
- Chance moves
- Bayesian games
- Fictitious play and stimulus response learning

2

An evolutionary perspective

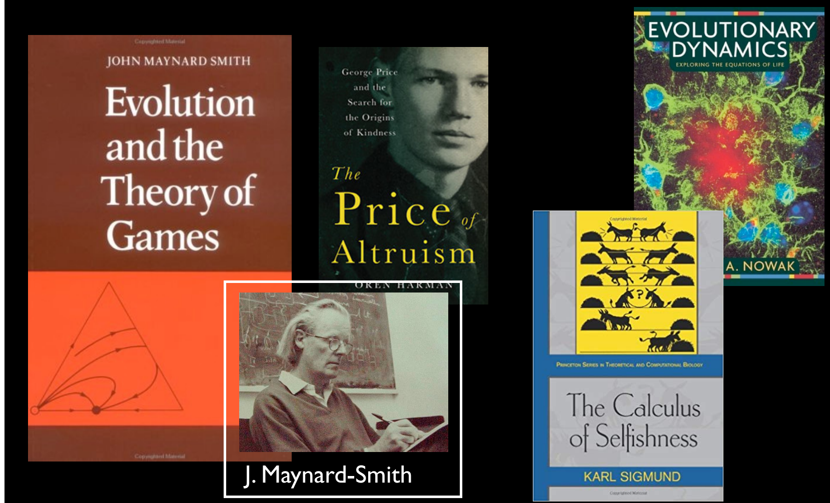

3

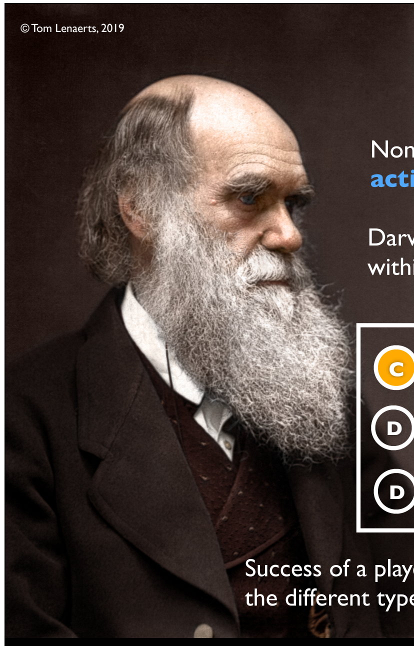

Non-rational players: **one player for each action**

Darwinian competition between simple players within **populations**

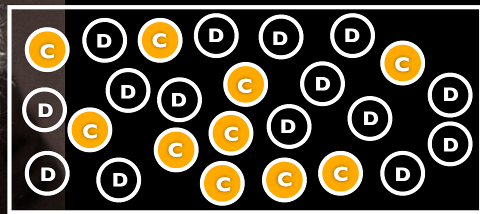

Success of a player **depends on the frequencies** of the different types of players

4

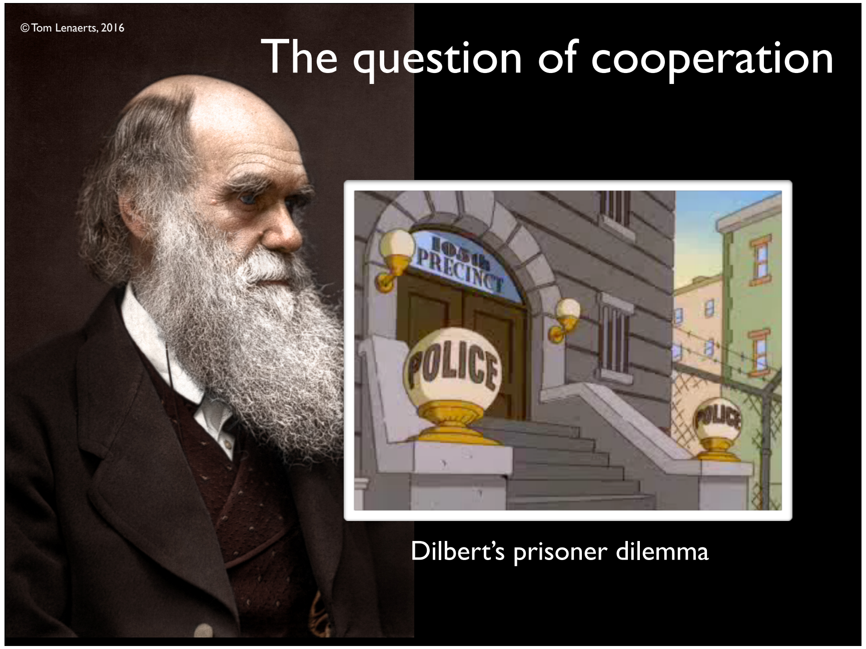
©Tom Lenaerts, 2016
The question of cooperation
105th PRECINCT
POLICE
Dilbert's prisoner dilemma

5-1

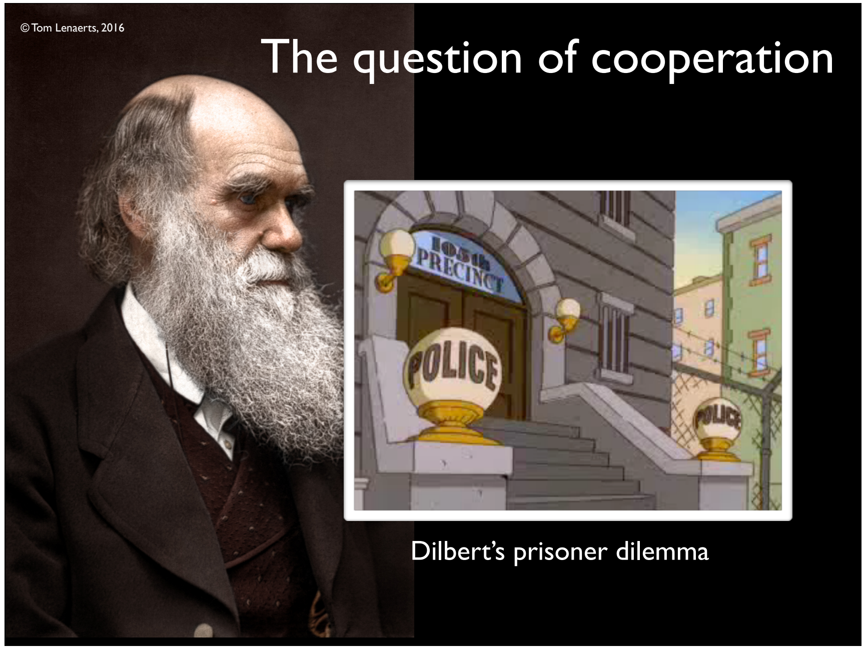
©Tom Lenaerts, 2016
The question of cooperation
105th PRECINCT
POLICE
Dilbert's prisoner dilemma

5-2


Social dilemmas

6

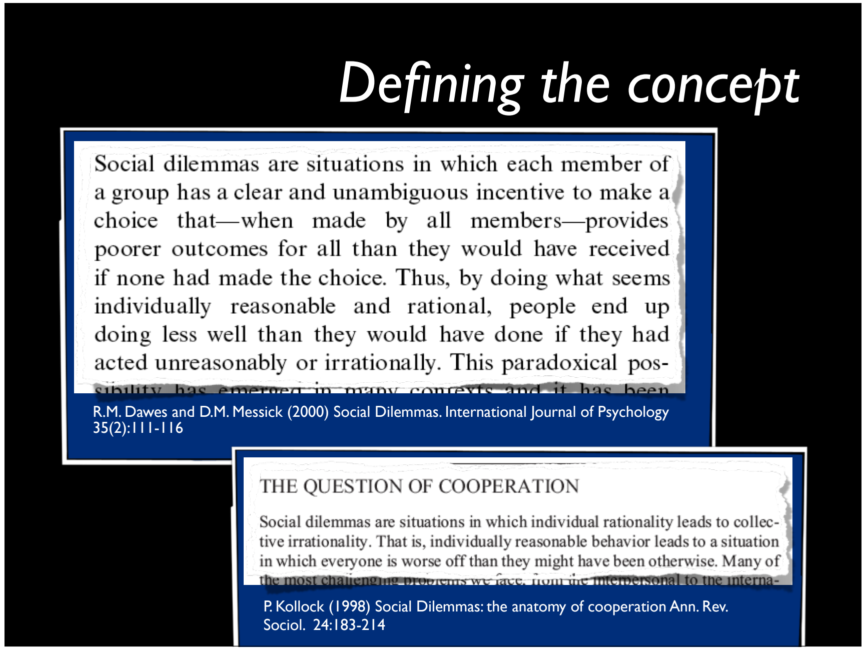
Defining the concept
Social dilemmas are situations in which each member of a group has a clear and unambiguous incentive to make a choice that—when made by all members—provides poorer outcomes for all than they would have received if none had made the choice. Thus, by doing what seems individually reasonable and rational, people end up doing less well than they would have done if they had acted unreasonably or irrationally. This paradoxical pos-
R.M. Dawes and D.M. Messick (2000) Social Dilemmas. International Journal of Psychology 35(2):111-116
THE QUESTION OF COOPERATION
Social dilemmas are situations in which individual rationality leads to collective irrationality. That is, individually reasonable behavior leads to a situation in which everyone is worse off than they might have been otherwise. Many of
P. Kollock (1998) Social Dilemmas: the anatomy of cooperation Ann. Rev. Sociol. 24:183-214

7

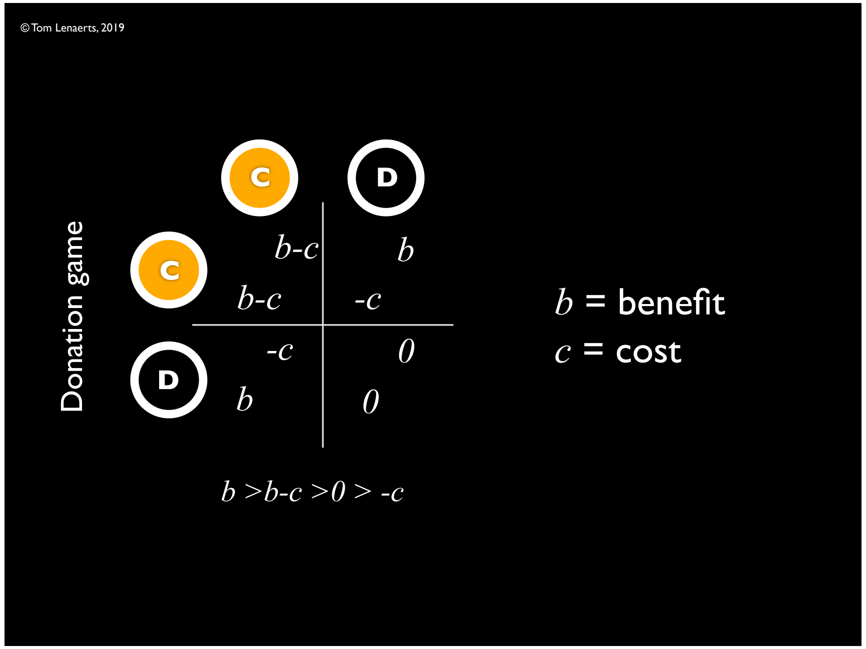
©Tom Lenaerts, 2019
Donation game
C
D
C
D
b-c
b
b-c
-c
-c
0
b
0
b = benefit
c = cost
b >b-c >0 > -c

8

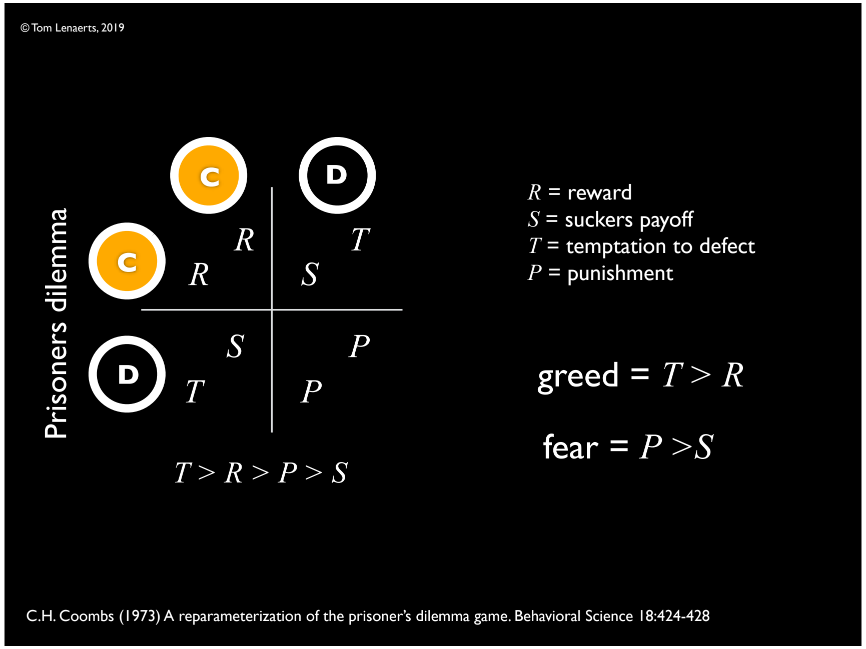
©Tom Lenaerts, 2019
Prisoners dilemma
C
D
C
D
R
T
R
S
S
P
T
P
T > R > P > S
R = reward
S = suckers payoff
T = temptation to defect
P = punishment
greed = T > R
fear = P >S
C.H. Coombs (1973) A reparameterization of the prisoner's dilemma game. Behavioral Science 18:424-428

9

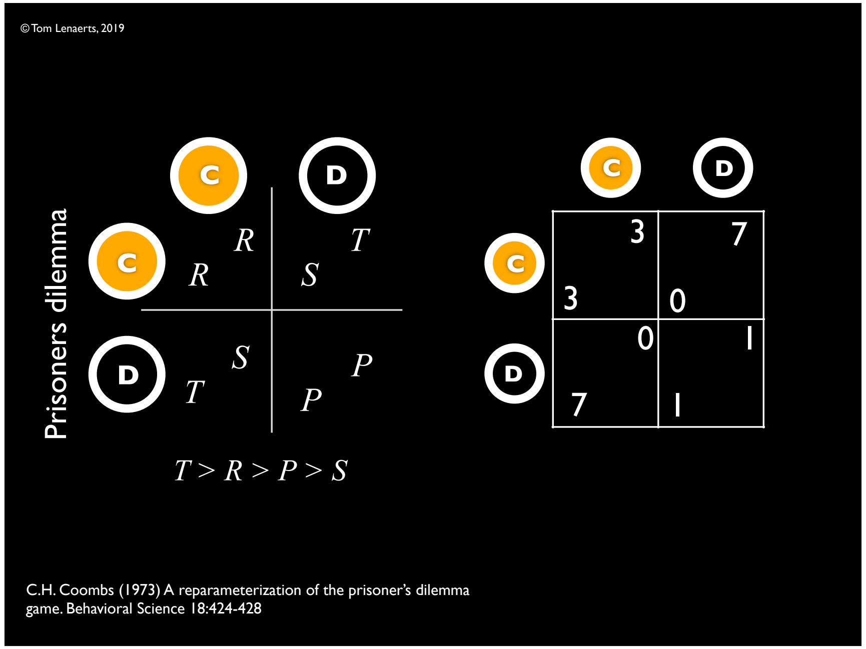
©Tom Lenaerts, 2019
Prisoners dilemma
C
D
C
D
R
T
R
S
S
P
T
P
T > R > P > S
3
7
3
0
0
1
7
1
C.H. Coombs (1973) A reparameterization of the prisoner's dilemma game. Behavioral Science 18:424-428

10-1

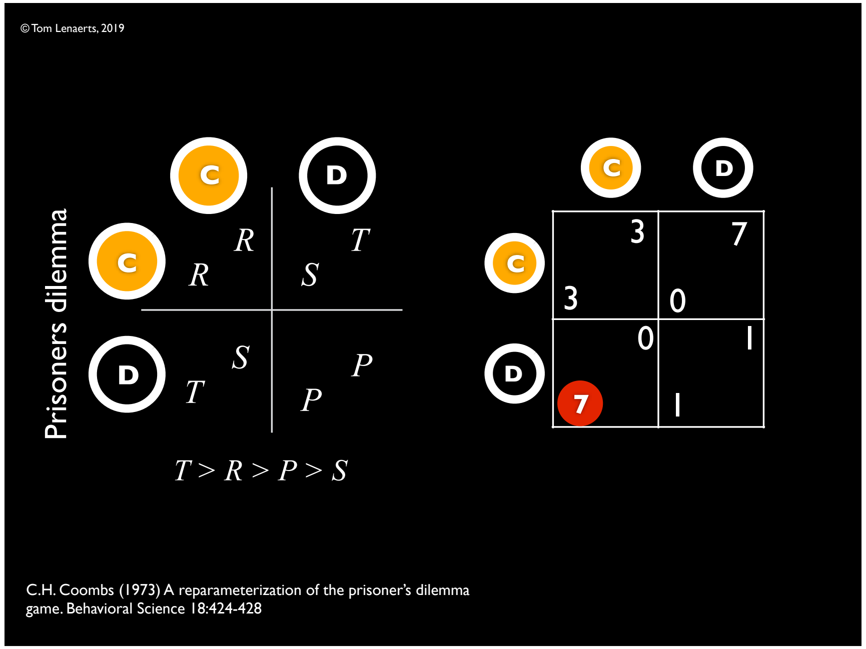
©Tom Lenaerts, 2019
Prisoners dilemma
C
D
C
D
R
T
R
S
S
P
T
P
T > R > P > S
3
7
3
0
0
1
7
1
C.H. Coombs (1973) A reparameterization of the prisoner's dilemma game. Behavioral Science 18:424-428

10-2

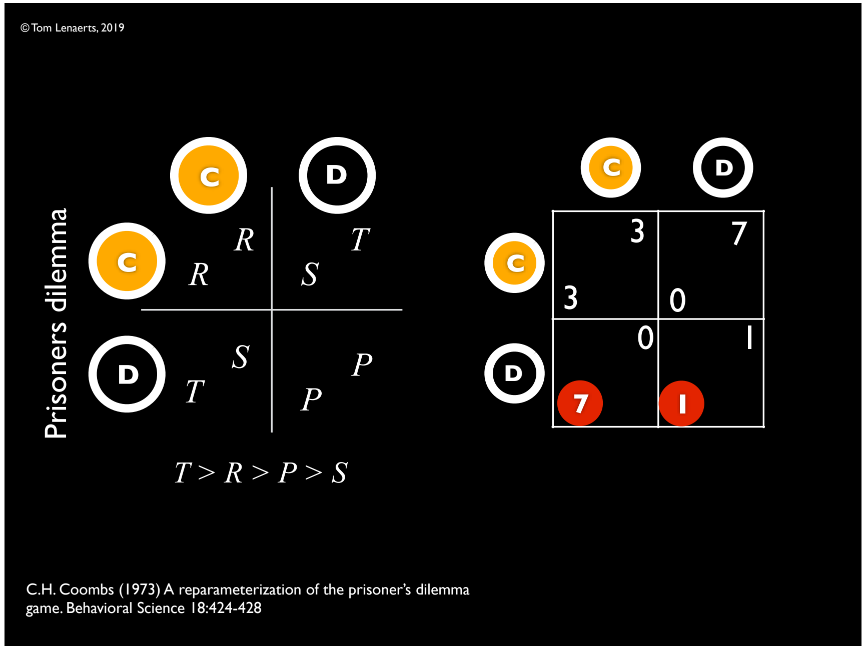

10-3

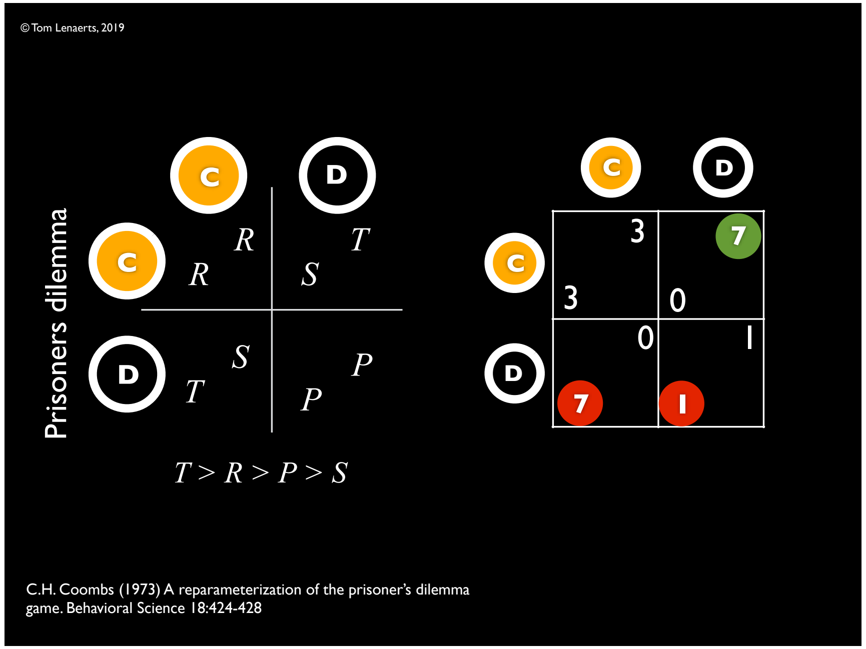

10-4

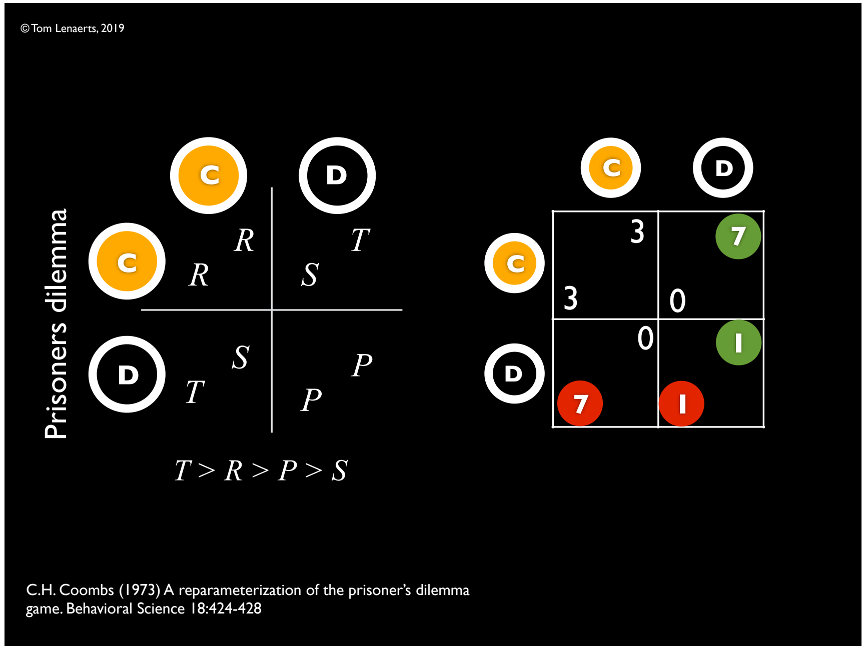

10-5

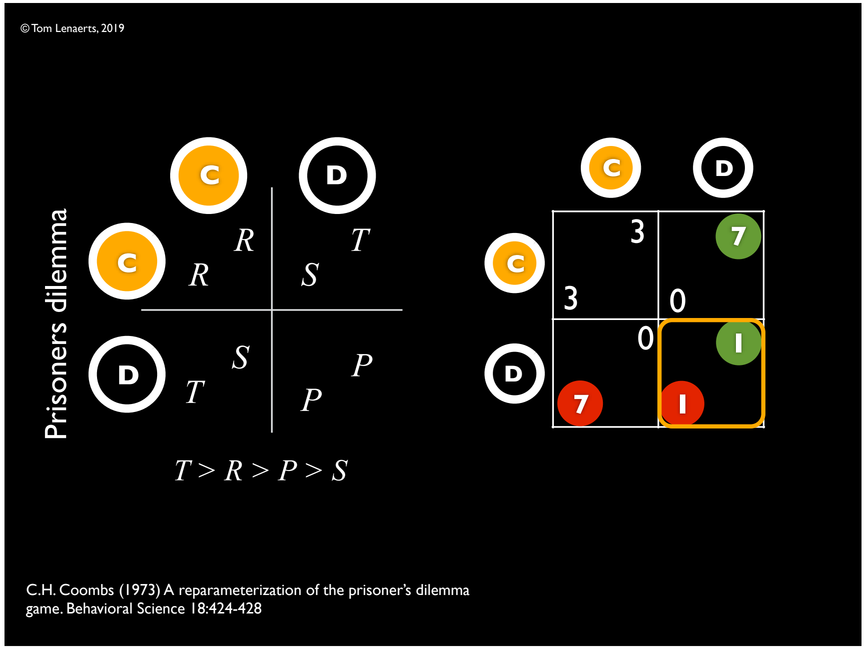

10-6

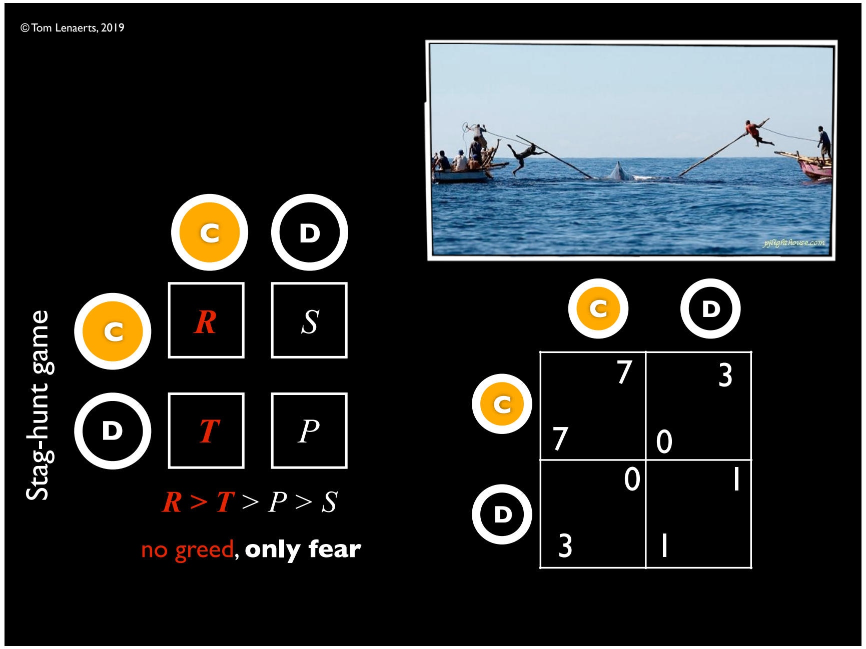

11-1

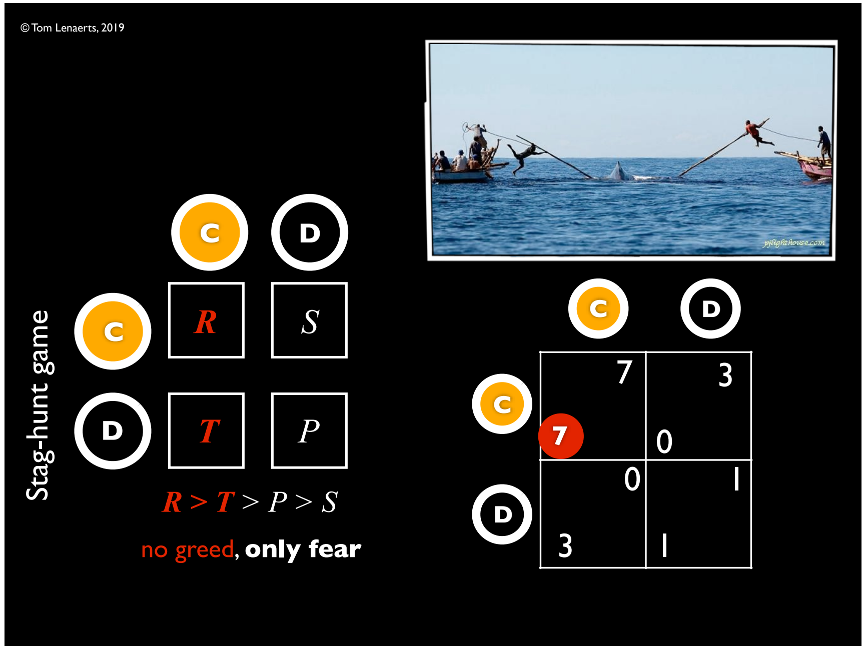

11-2

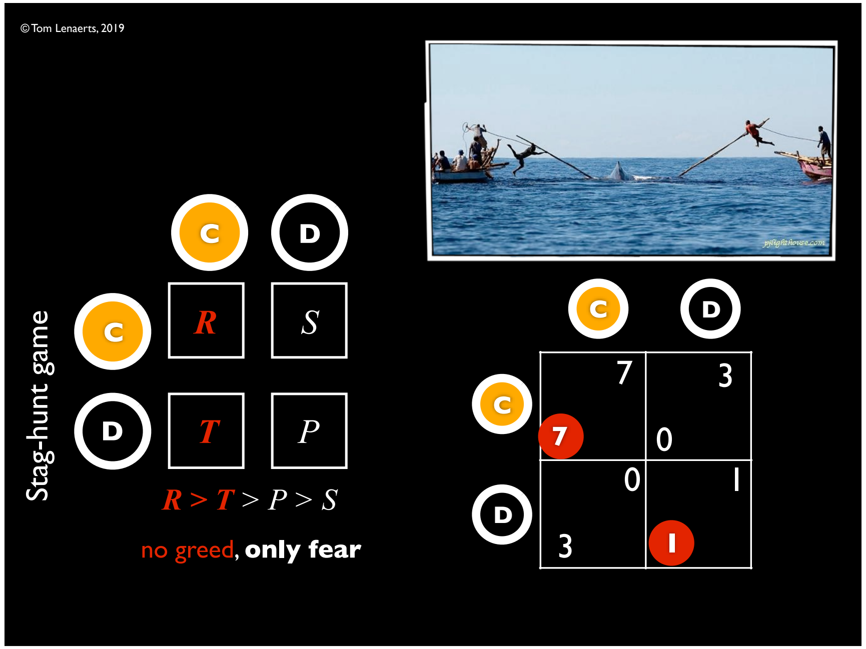

11-3

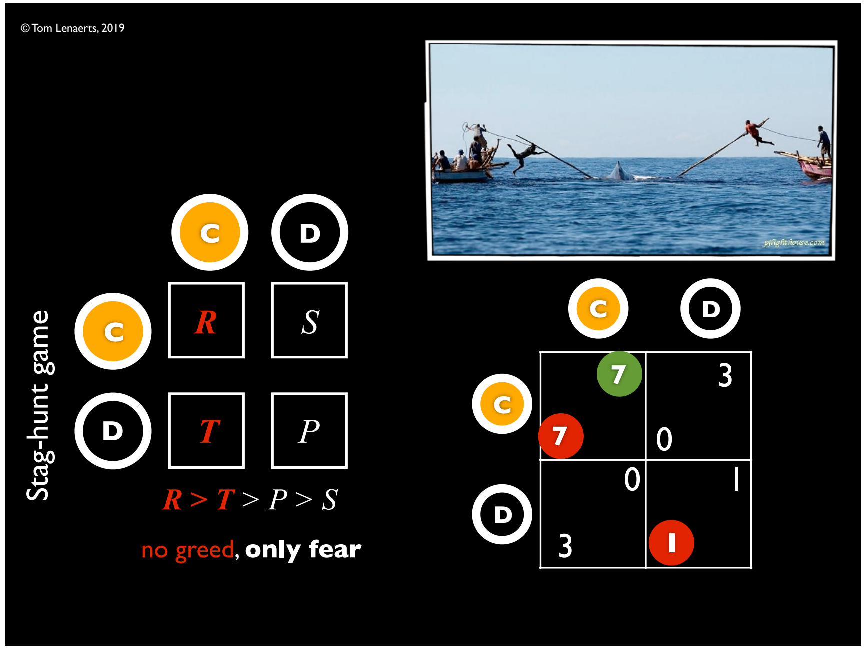

11-4

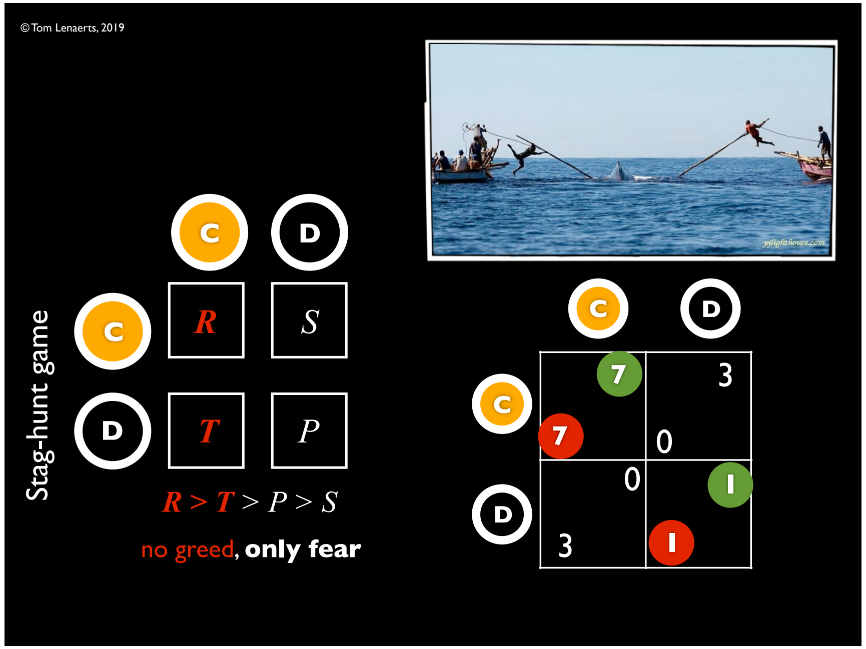

11-5

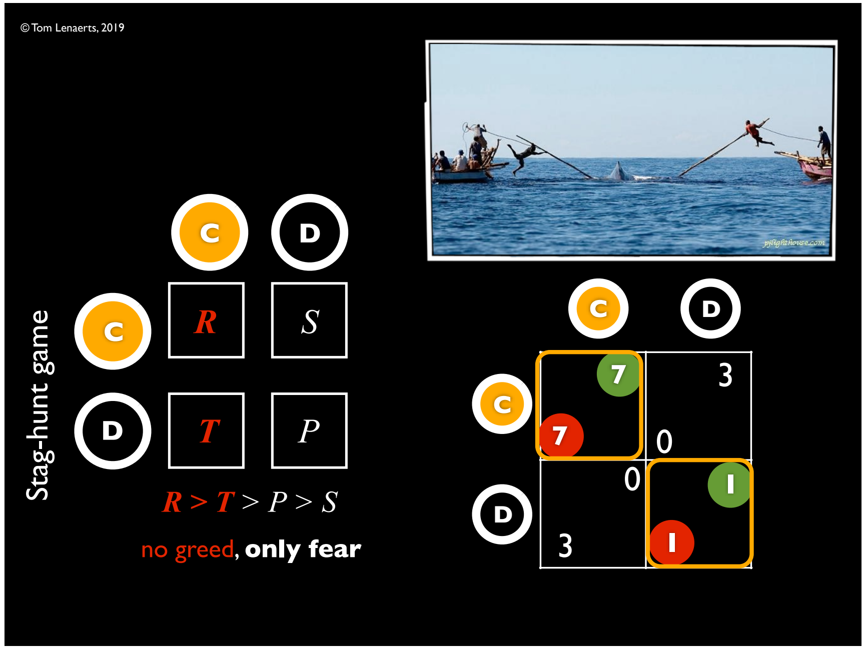

11-6

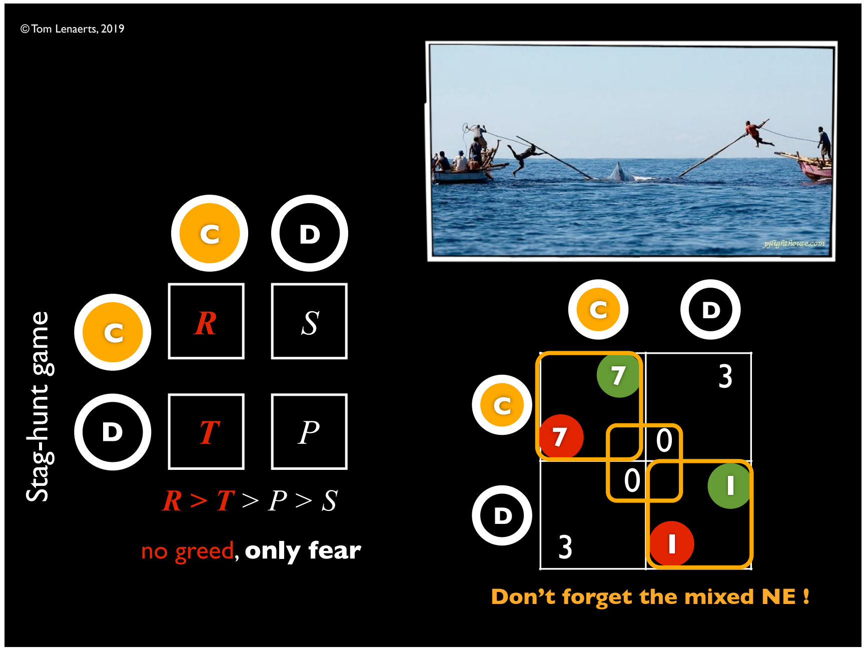

11-7

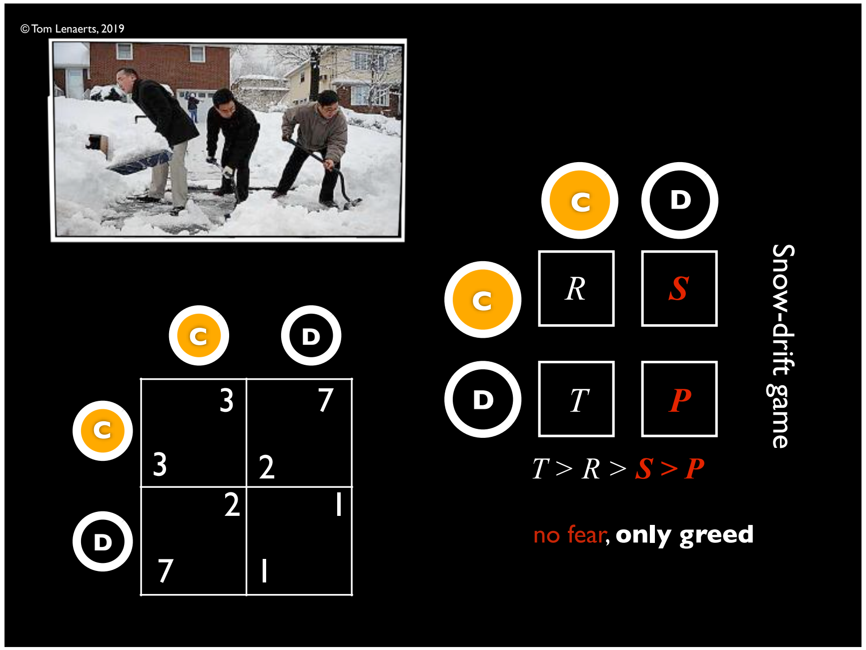

12-1

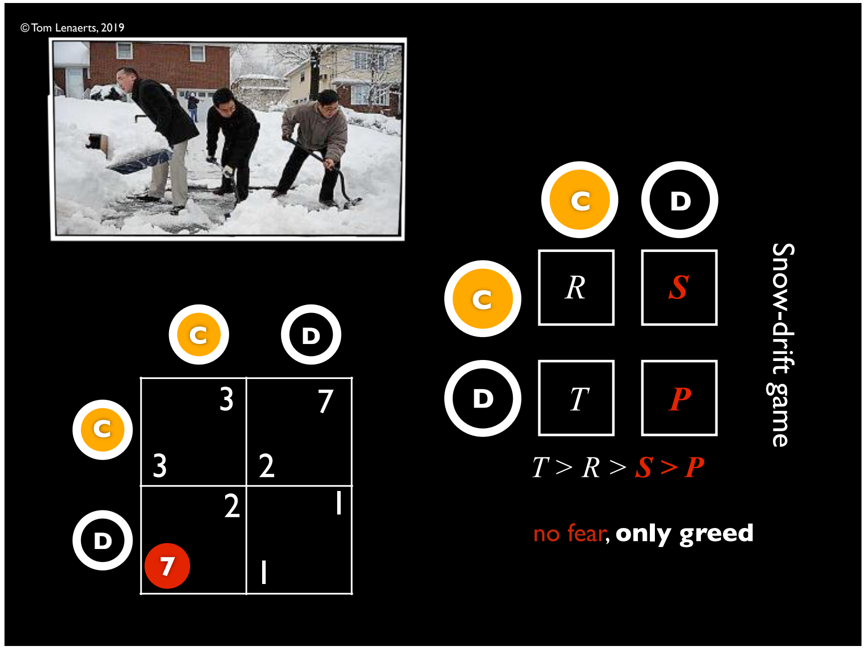

12-2

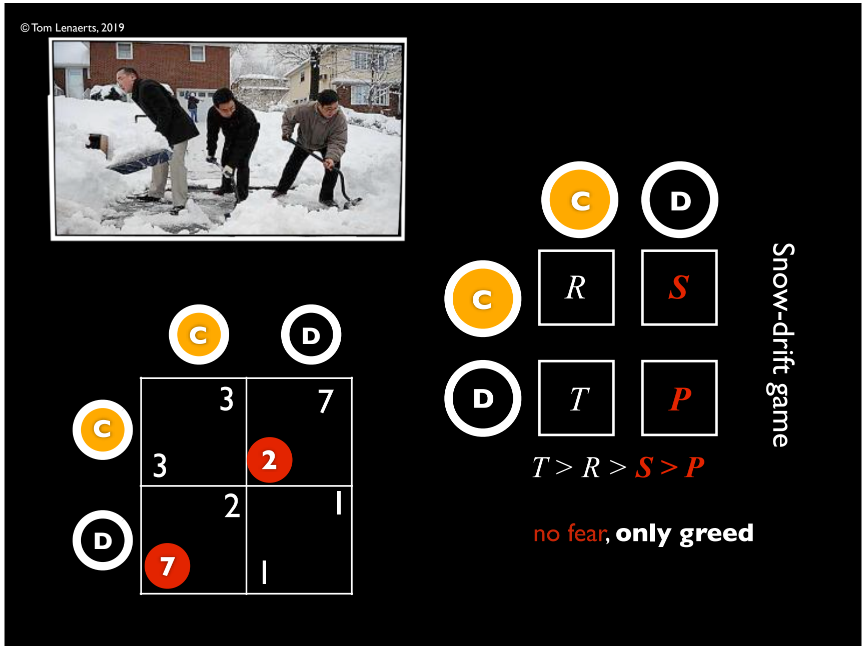

12-3

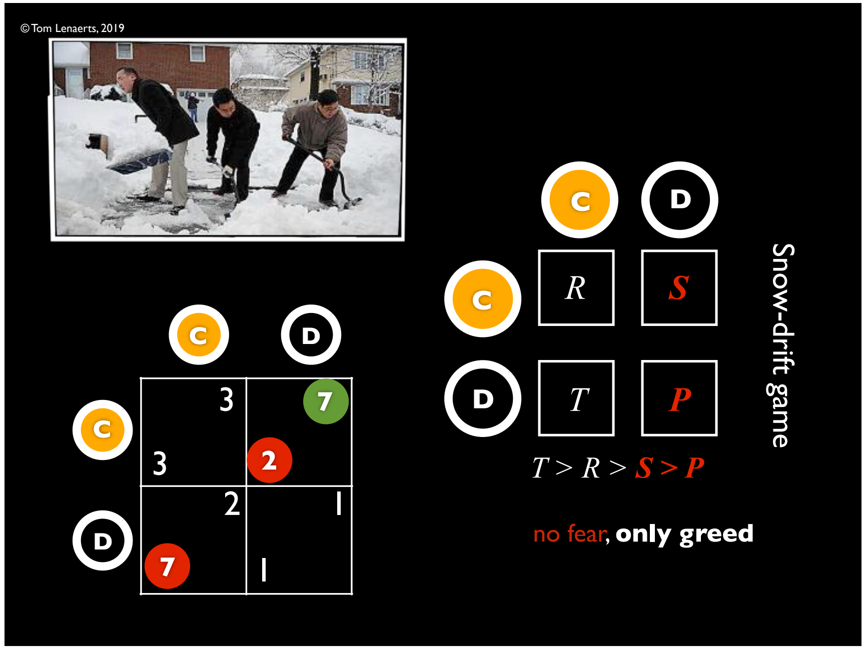

12-4

12-5

12-6

12-7

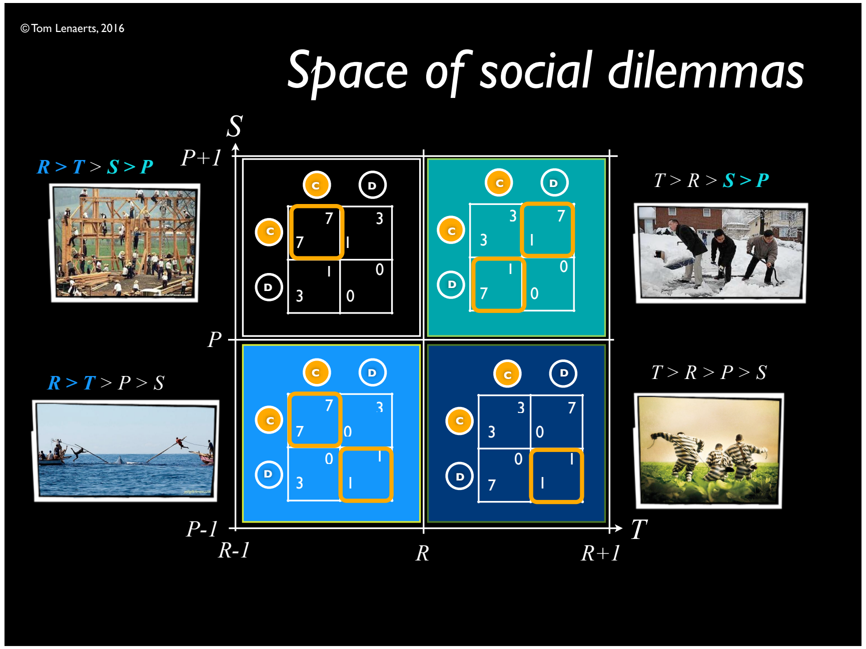

13

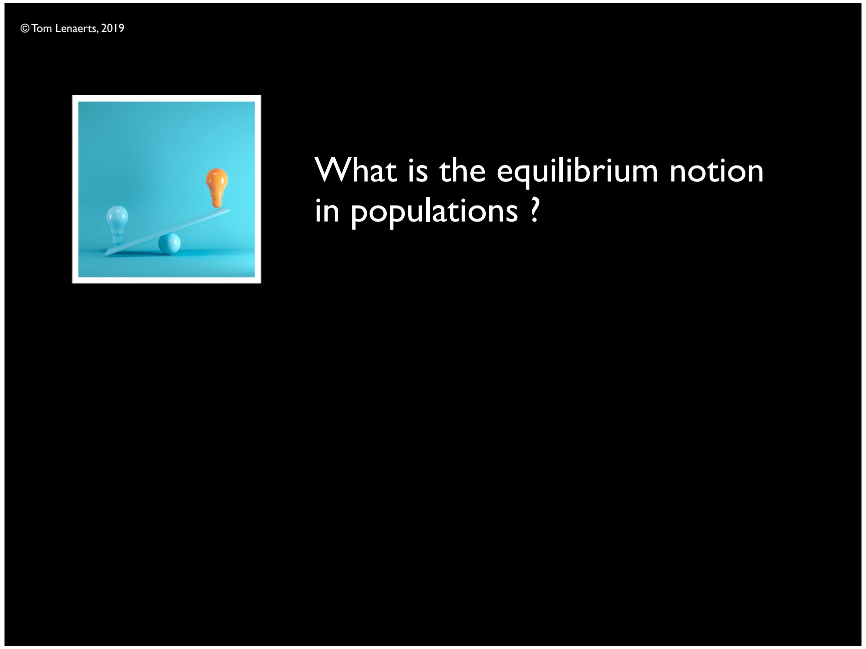

14-1

What is the equilibrium notion in populations ?

John Maynard Smith and George Price (1973)

The evolutionary stable strategy (ESS) concept

14-2

A strategy is an **ESS** when it cannot be invaded by another strategy

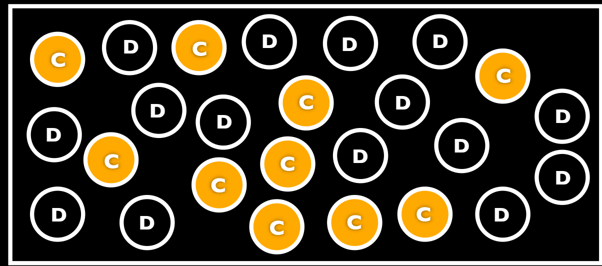

Infinite population assumption

J. Maynard-Smith and G.R. Price (1973) The logic of animal conflict. Nature 246:15-18

15

Can **C** invade a population of **D** players

The fraction of **C** (**D**) players is ε $(1-\varepsilon)$

success of **C** in a **D** population | success of **D** in a **D** population

$$S(1-\varepsilon)+R\varepsilon > P(1-\varepsilon)+T\varepsilon$$

	C	D
C	R	S
D	T	P

C can invade when:

If **C** cannot invade then **D** is an **ESS**

J. Maynard-Smith and G.R. Price (1973) The logic of animal conflict. Nature 246:15-18

16-1

Can **C** invade a population of **D** players

The fraction of **C** (**D**) players is ε $(1-\varepsilon)$

success of **C** in a **D** population | success of **D** in a **D** population

$$S(1-\varepsilon)+R\varepsilon > P(1-\varepsilon)+T\varepsilon$$

	C	D
C	R	S
D	T	P

C can invade when:

i) $S>P$

If **C** cannot invade then **D** is an **ESS**

J. Maynard-Smith and G.R. Price (1973) The logic of animal conflict. Nature 246:15-18

16-2

Can **C** invade a population of **D** players

The fraction of **C** (**D**) players is ε $(1\text{-}\varepsilon)$

success of **C** in a **D** population | success of **D** in a **D** population

$$S(1-\varepsilon)+R\varepsilon > P(1-\varepsilon)+T\varepsilon$$

	C	D
C	R	S
D	T	P

C can invade when:

i) $S>P$

ii) $S=P$ and $R>T$

If **C** cannot invade then **D** is an **ESS**

J. Maynard-Smith and G.R. Price (1973) The logic of animal conflict. Nature 246:15-18

16-3

Can **D** invade a population of **C** players

The fraction of **D** (**C**) players is ε $(1\text{-}\varepsilon)$

success of **D** in a **C** population | success of **C** in a **C** population

$$T(1-\varepsilon)+P\varepsilon > R(1-\varepsilon)+S\varepsilon$$

	C	D
C	R	S
D	T	P

D can invade when:

If **D** cannot invade then **C** is an **ESS**

J. Maynard-Smith and G.R. Price (1973) The logic of animal conflict. Nature 246:15-18

17-1

Can **D** invade a population of **C** players

The fraction of **D** (**C**) players is ε $(1\text{-}\varepsilon)$

success of **D** in a **C** population | success of **C** in a **C** population

$$T(1-\varepsilon)+P\varepsilon > R(1-\varepsilon)+S\varepsilon$$

	C	D
C	R	S
D	T	P

D can invade when:

i) $T>R$

If **D** cannot invade then **C** is an **ESS**

J. Maynard-Smith and G.R. Price (1973) The logic of animal conflict. Nature 246:15-18

17-2

Can **D** invade a population of **C** players

The fraction of **D** (**C**) players is ε $(1\text{-}\varepsilon)$

success of **D** in a **C** population | success of **C** in a **C** population

$$T(1-\varepsilon)+P\varepsilon > R(1-\varepsilon)+S\varepsilon$$

	C	D
C	R	S
D	T	P

D can invade when:

i) $T>R$

ii) $T=R$ and $P>S$

If **D** cannot invade then **C** is an **ESS**

J. Maynard-Smith and G.R. Price (1973) The logic of animal conflict. Nature 246:15-18

17-3

Can **C** invade **D**?



	C	D
C	3, 3	0, 7
D	7, 0	1, 1

	C	D
C	7, 7	0, 3
D	3, 0	1, 1

	C	D
C	3, 3	2, 7
D	7, 2	1, 1

18-1

Can **C** invade **D**?



	C	D
C	3, 3	0, 7
D	7, 0	1, 1

	C	D
C	7, 7	0, 3
D	3, 0	1, 1

	C	D
C	3, 3	2, 7
D	7, 2	1, 1

C no since $P>S$

18-2

Can **C** invade **D**?



	C	D
C	3, 3	0, 7
D	7, 0	1, 1

	C	D
C	7, 7	0, 3
D	3, 0	1, 1

	C	D
C	3, 3	2, 7
D	7, 2	1, 1

C no since $P>S$

C no since $P>S$

18-3

Can **C** invade **D**?



	C	D
C	3, 3	0, 7
D	7, 0	1, 1

	C	D
C	7, 7	0, 3
D	3, 0	1, 1

	C	D
C	3, 3	2, 7
D	7, 2	1, 1

C no since $P>S$

C no since $P>S$

C yes since $P<S$

18-4

Can D invade C?

	C	D
C	3, 3	0, 7
D	7, 0	1, 1

C: no since $P>S$

	C	D
C	7, 7	0, 3
D	3, 0	1, 1

C: no since $P>S$

	C	D
C	3, 3	2, 7
D	7, 2	1, 1

C: yes since $P<S$

18-5

Can D invade C?

	C	D
C	3, 3	0, 7
D	7, 0	1, 1

C: no since $P>S$

D: yes since $R<T$

	C	D
C	7, 7	0, 3
D	3, 0	1, 1

C: no since $P>S$

	C	D
C	3, 3	2, 7
D	7, 2	1, 1

C: yes since $P<S$

18-6

Can D invade C?

	C	D
C	3, 3	0, 7
D	7, 0	1, 1

C: no since $P>S$

D: yes since $R<T$

	C	D
C	7, 7	0, 3
D	3, 0	1, 1

C: no since $P>S$

D: no since $R>T$

	C	D
C	3, 3	2, 7
D	7, 2	1, 1

C: yes since $P<S$

18-7

Can D invade C?

	C	D
C	3, 3	0, 7
D	7, 0	1, 1

C: no since $P>S$

D: yes since $R<T$

	C	D
C	7, 7	0, 3
D	3, 0	1, 1

C: no since $P>S$

D: no since $R>T$

	C	D
C	3, 3	2, 7
D	7, 2	1, 1

C: yes since $P<S$

D: yes since $R<T$

18-8

Which one is an ESS?

	C	D
C	3, 3	0, 7
D	7, 0	1, 1

D Is an **ESS**

	C	D
C	7, 7	0, 3
D	3, 0	1, 1

C is an **ESS**

D is an **ESS**

	C	D
C	3, 3	2, 7
D	7, 2	1, 1

no **ESS** ?

19-1

Which one is an ESS?

	C	D
C	3, 3	0, 7
D	7, 0	1, 1

D Is an **ESS**

	C	D
C	7, 7	0, 3
D	3, 0	1, 1

C is an **ESS**

D is an **ESS**

	C	D
C	3, 3	2, 7
D	7, 2	1, 1

Don't forget the mixed NE !

no **ESS** ?

19-2

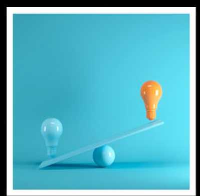

What is the equilibrium notion in populations ?

20-1

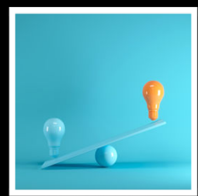

What is the equilibrium notion in populations ?

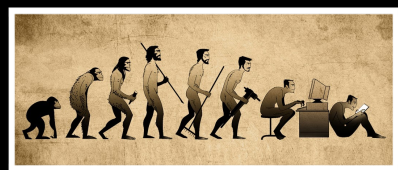

How does learning work in a population?

20-2

Change in populations

Evolutionary dynamics

Imitate successful behaviour (a.k.a. social learning)

Genetic evolution of successful properties (a.k.a. survival of the fittest)

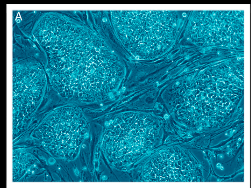

21

In *infinite* populations

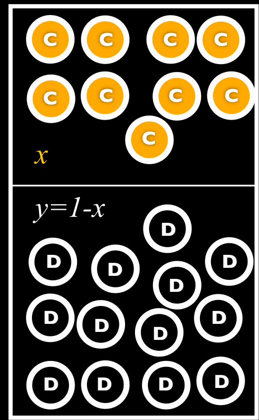

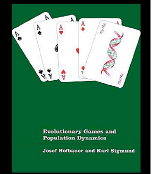

Replicator equation ...

$$\frac{dx}{dt} = x(1-x)[f_C(x)-f_D(x)]$$

$$= x(1-x)[(b-c+c-b+0)x-c-0]$$

$$= -cx(1-x)$$

C ⟶ D

$x=1$ $x=0$

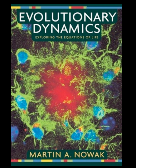

P.D. Taylor and L.B. Jonker (1978) Evolutionary stable strategies and game dynamics. Mathematical biosciences 40(1-2):145-156

22

Change in *infinite* populations

$$\frac{dx}{dt} = x(1-x)[(R-S-T+P)x+S-P]$$

100% C —— **D dominates C** ⟶ 100% D: prisoners dilemma

100% C ⟵ x^* ⟶ 100% D — **D and C are bistable**: Stag-hunt game

100% C ⟶ x^* ⟵ 100% D — **D and C coexist**: Snowdrift game

$$x^* = \frac{S-P}{R-S-T+P}$$

23

In all social dilemmas

$$\frac{S-P}{R-S-T+P}$$

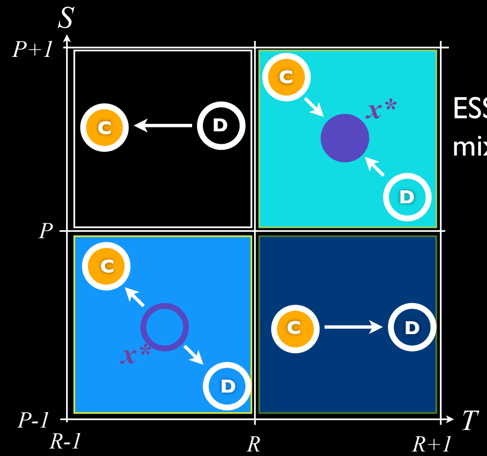

ESS = mixed state

24

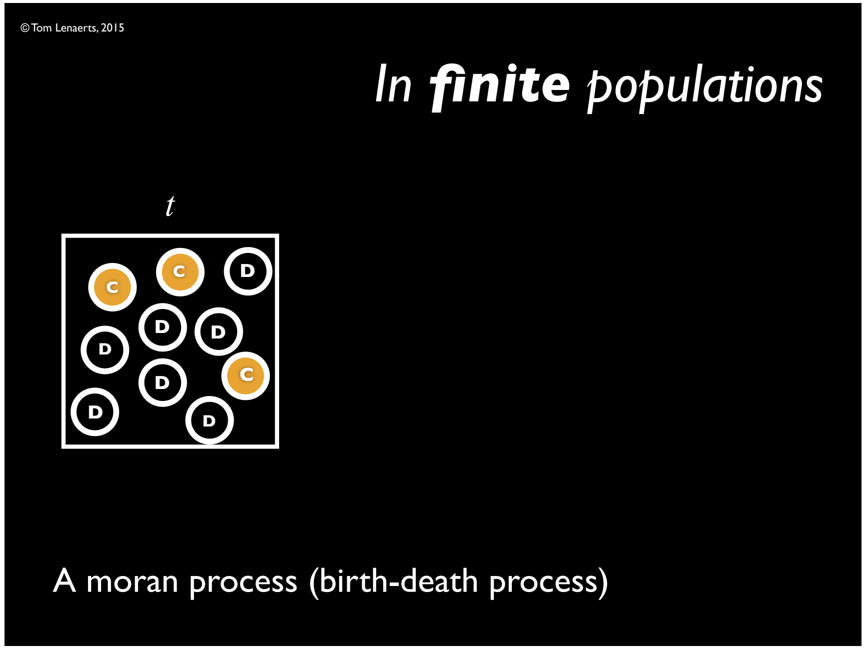

25-1

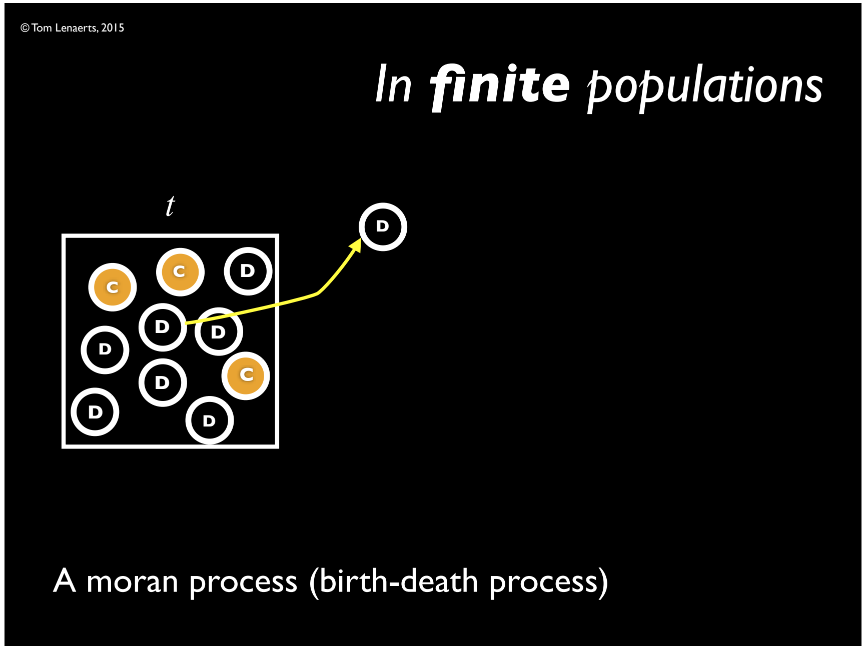

25-2

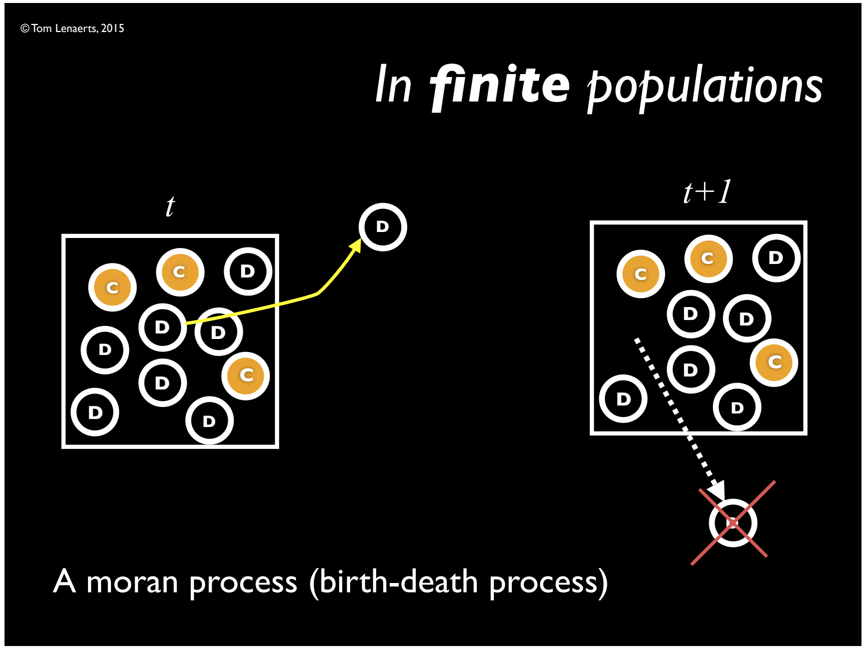

25-3

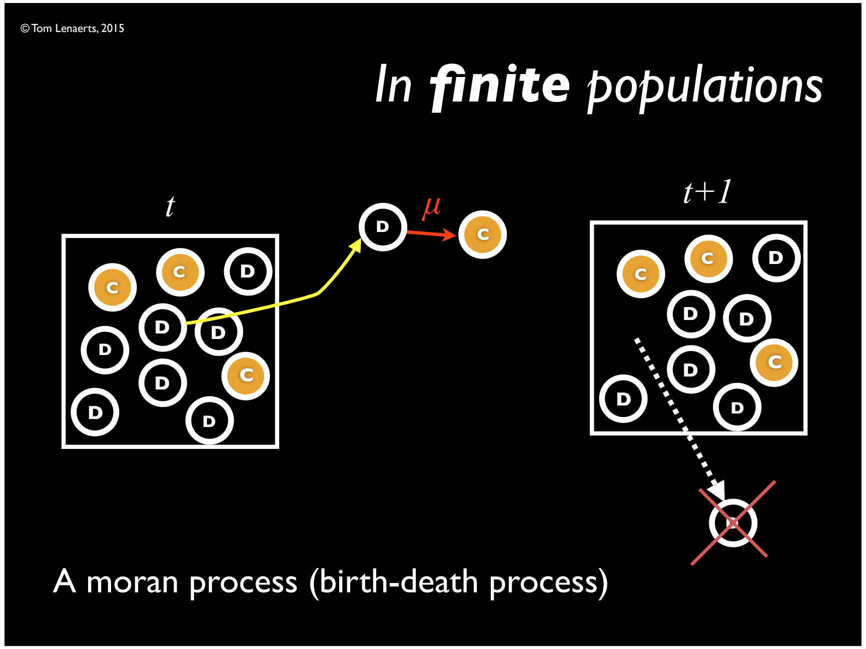

25-4

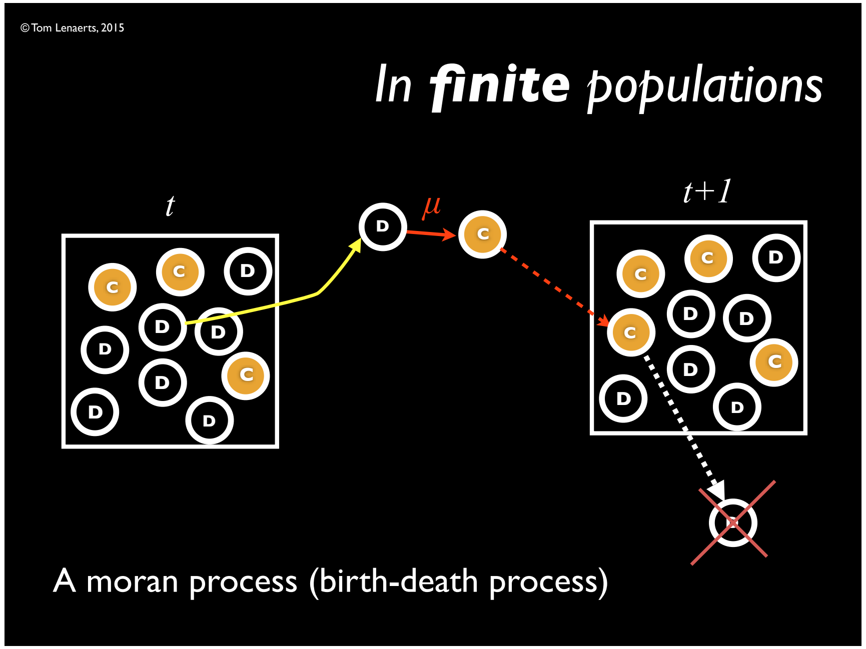

25-5

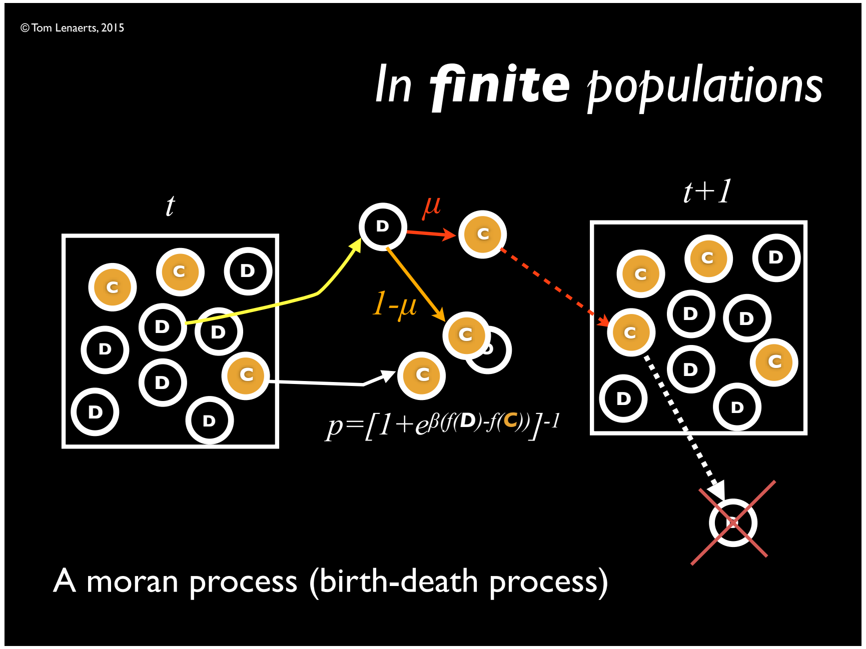

25-6

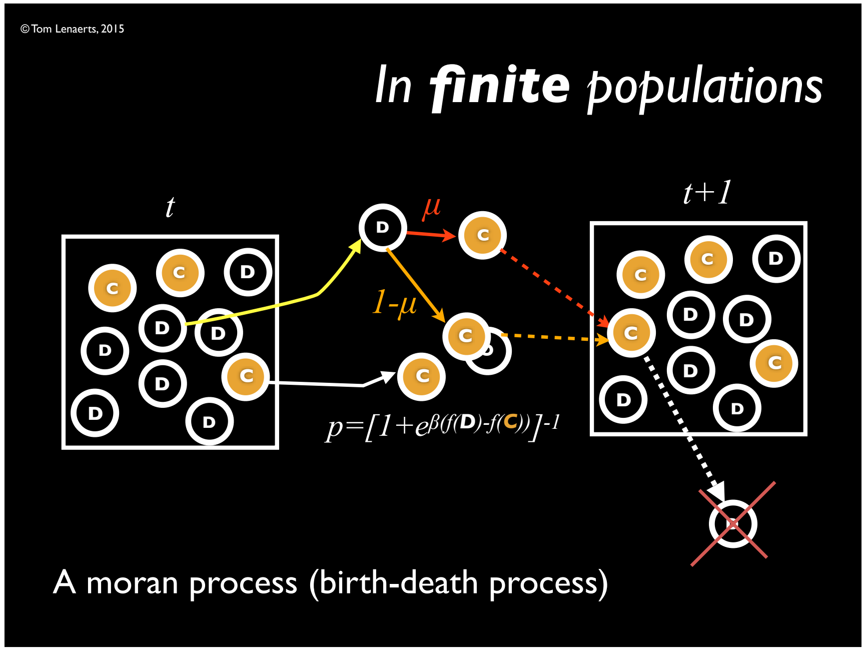

25-7

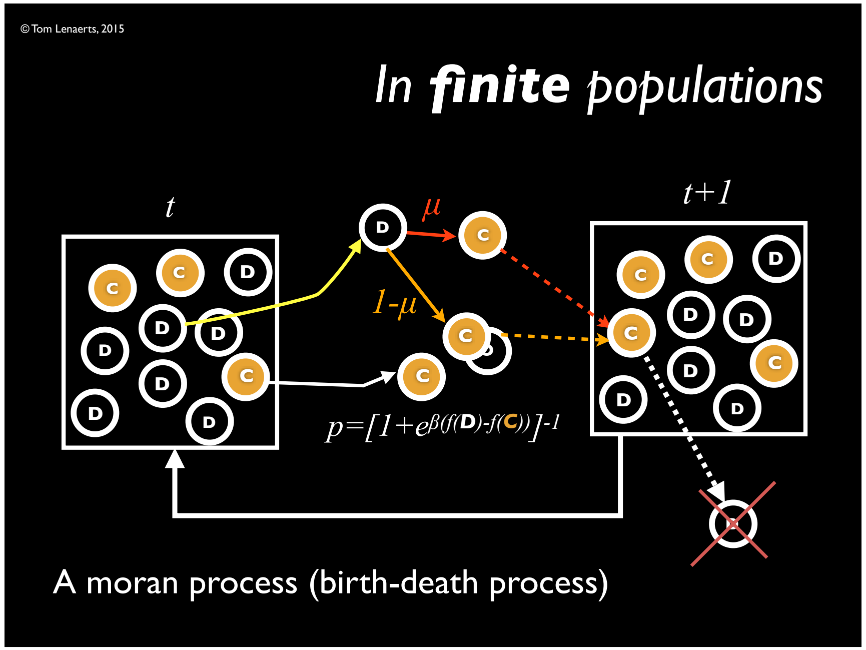

25-8

*In **finite** populations*

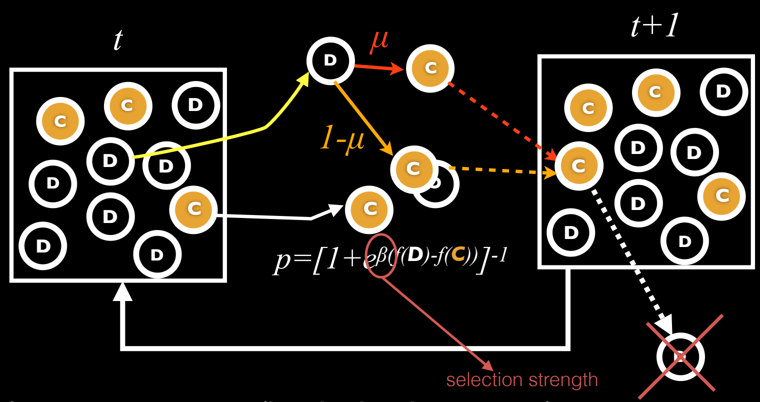

A moran process (birth-death process)

25-9

*In **finite** populations*

A moran step algorithm

26-1

*In **finite** populations*

A moran step algorithm

First, select randomly two players (with replacement)

26-2

*In **finite** populations*

A moran step algorithm

First, select randomly two players (with replacement)

Second, let each player play the game against all other players (not themselves).

26-3

*In **finite** populations*

A moran step algorithm

- First, select randomly two players (with replacement)
- Second, let each player play the game against all other players (not themselves).
- Third, calculate the average fitness of the player

26-4

*In **finite** populations*

A moran step algorithm

- First, select randomly two players (with replacement)
- Second, let each player play the game against all other players (not themselves).
- Third, calculate the average fitness of the player
- **If** a random value is smaller than the fermi probability

26-5

*In **finite** populations*

A moran step algorithm

- First, select randomly two players (with replacement)
- Second, let each player play the game against all other players (not themselves).
- Third, calculate the average fitness of the player
- **If** a random value is smaller than the fermi probability

Then first player ← second player **or** with probability μ the first ← random strategy
Else same but second ← first

26-6

In pseudo python

A moran step algorithm (without mutation):

```
def moran-step(beta, population):
    '''
    This function implements a birth-death process over
    the population. At time t, two players are randomly
    selected from the population
    '''
    selected=select_random_with_replacement(population, 2)
    for i, player in enumerate(selected):
        for j in range(len(population)):
            if j == player: continue
            players_payoffs = play_game(population[player],
                                        population[j])
            fitness[i] += players_payoffs[0]
    fitness = fitness / (Z-1)
    if random() < prob_imitation(beta, fitness):
        population[selected[0]] = population[selected[1]]
    else:
        population[selected[1]] = population[selected[0]]

    return population
```

27

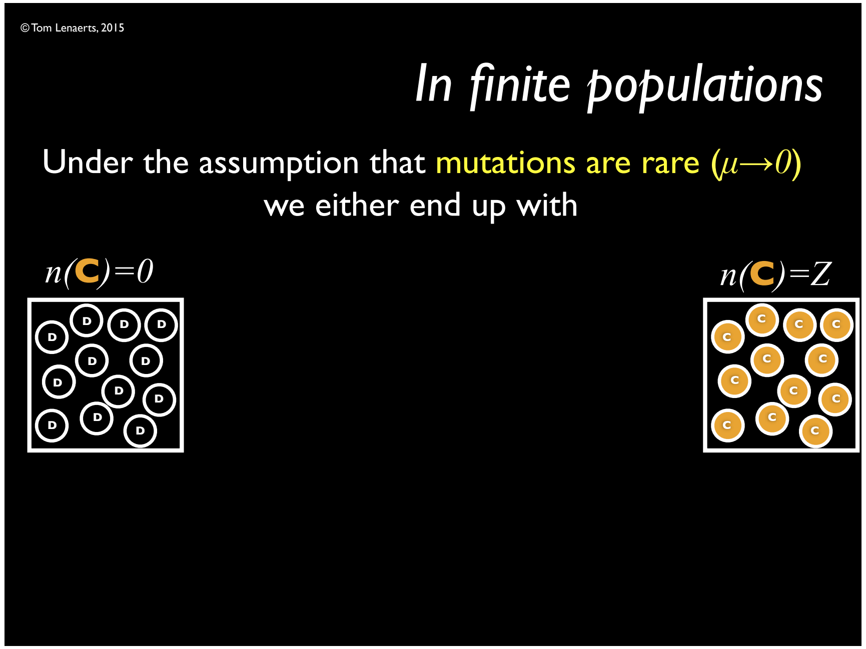

28-1

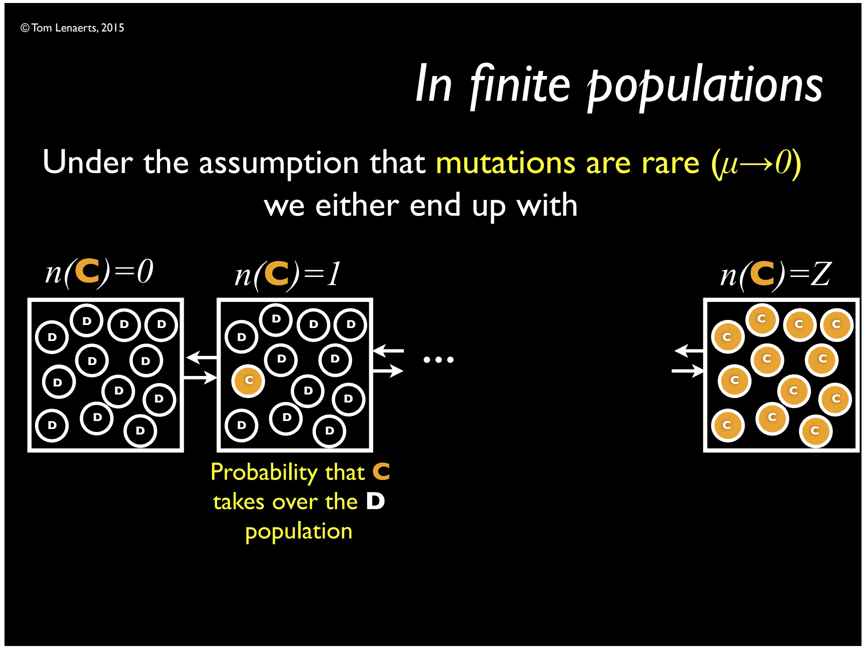

28-2

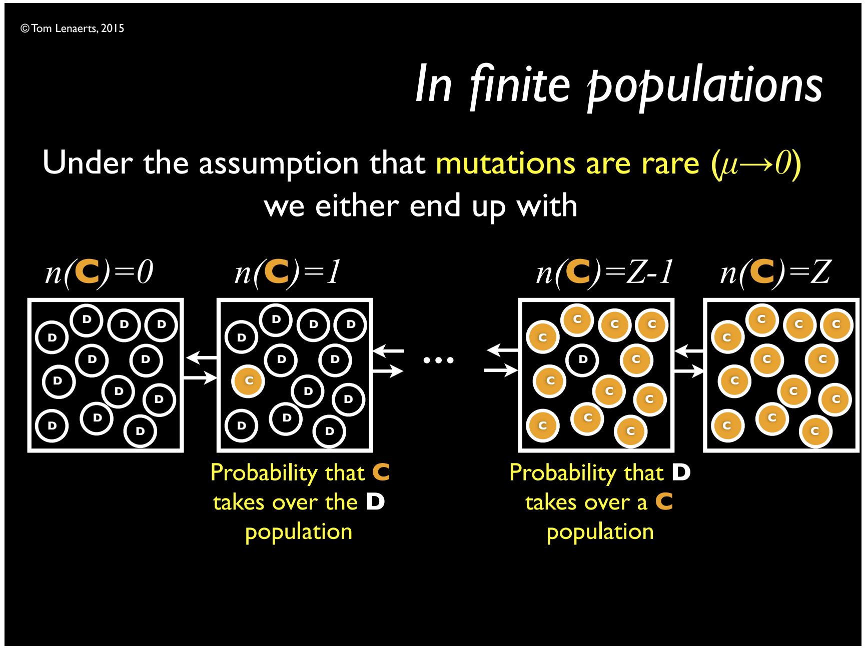

28-3

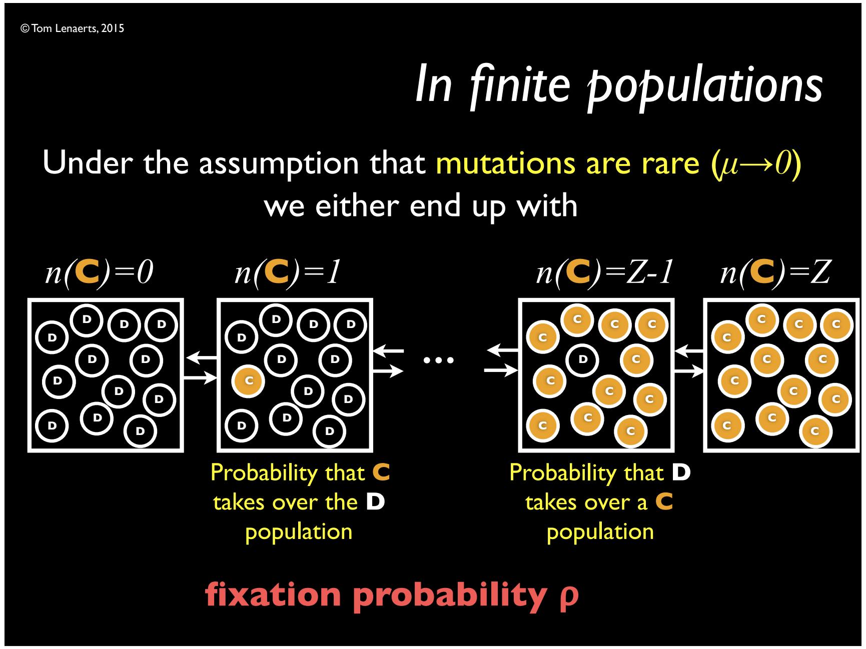

28-4

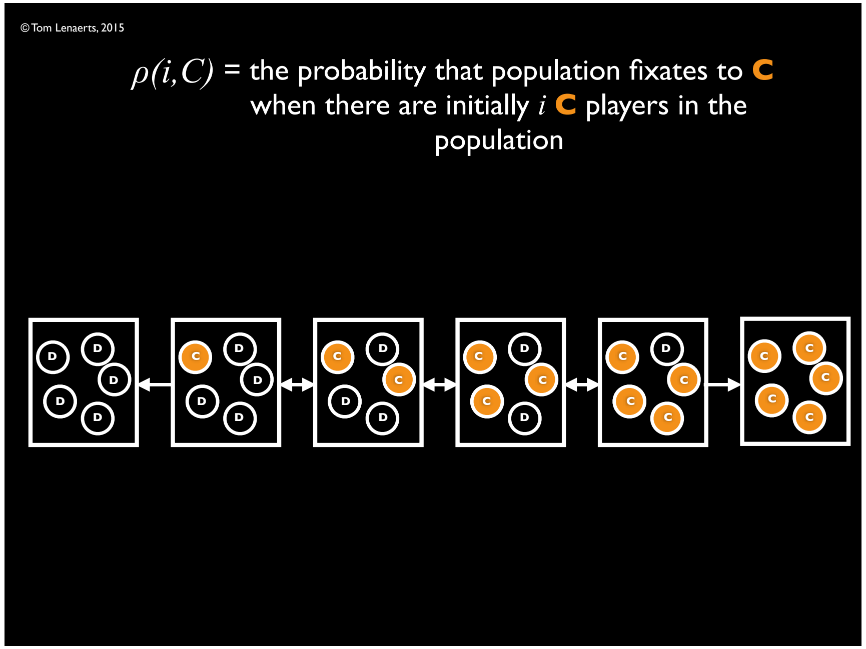

29-1

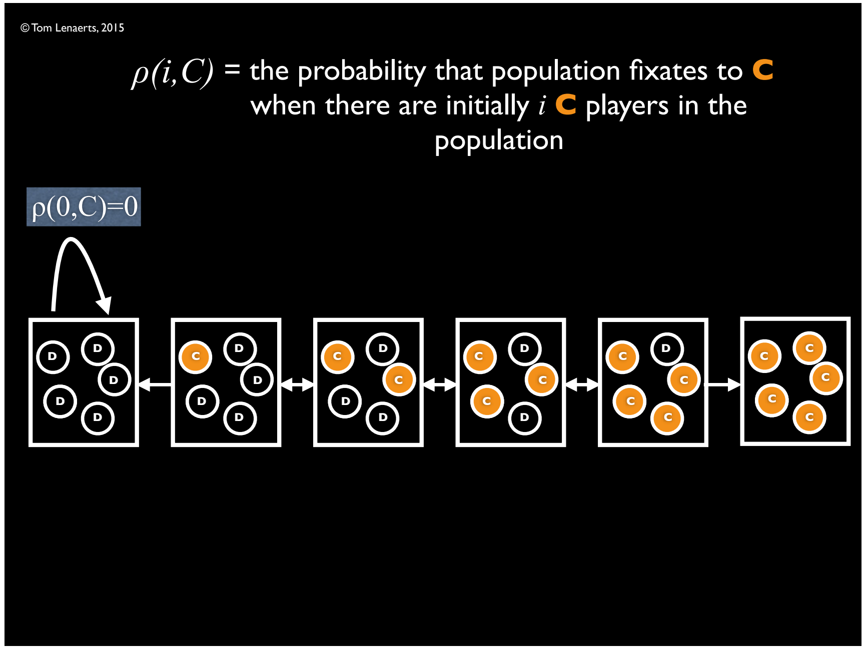

29-2

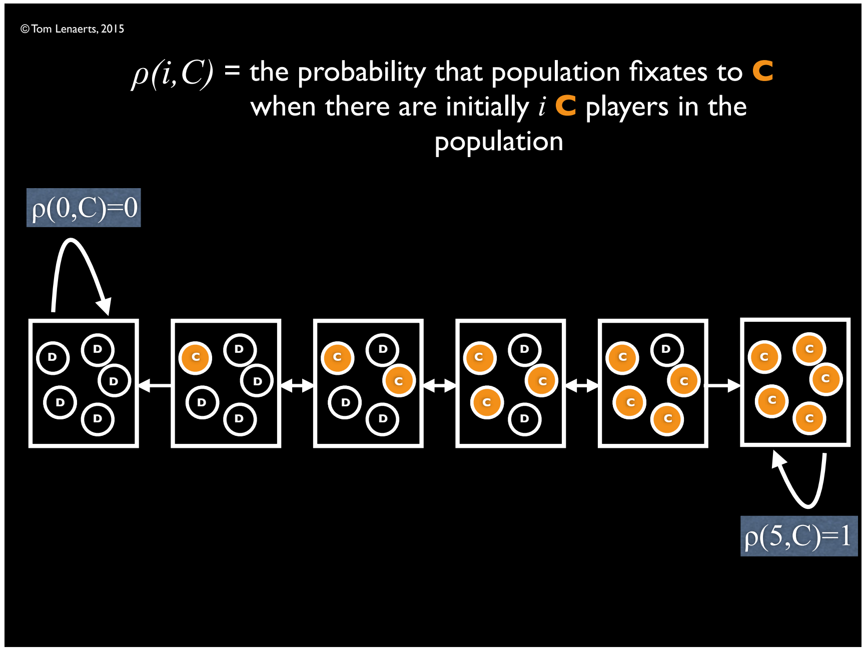

29-3

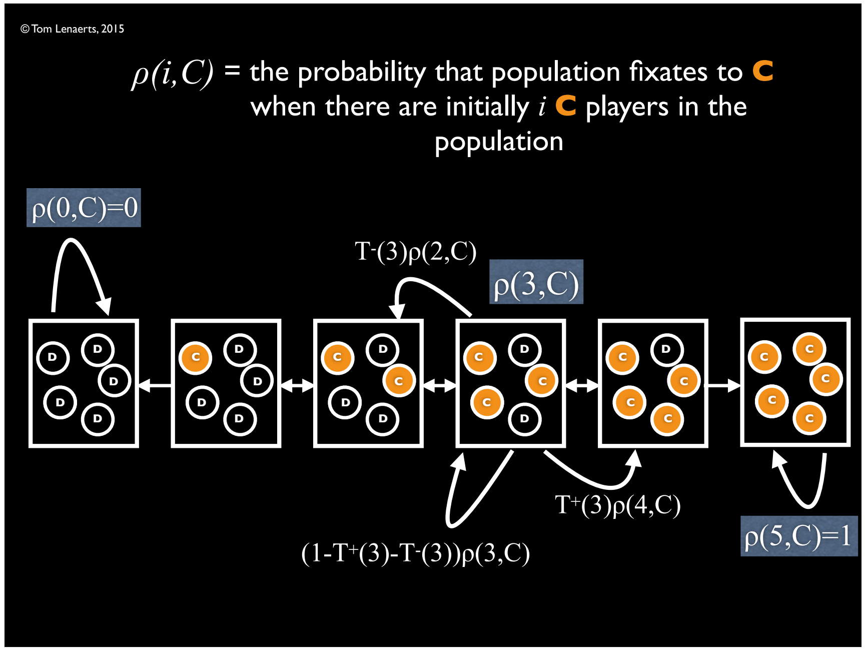

29-4

$\rho(i,C)$ = the probability that population fixates to **C** when there are initially i **C** players in the population

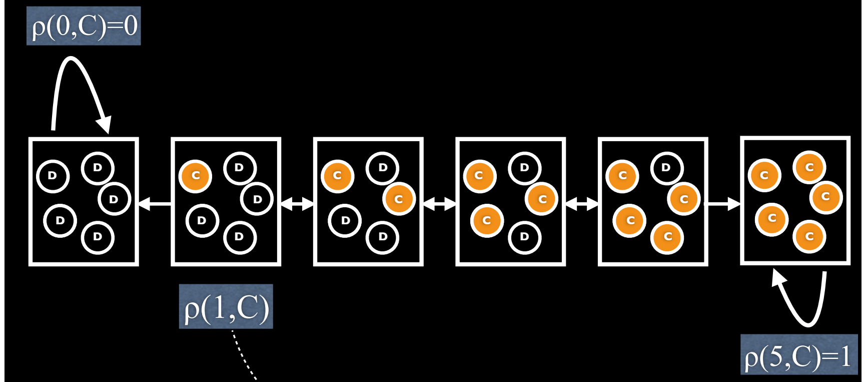

$\rho(1,C) = T^-(1)\rho(0,C) + T^+(1)\rho(2,C) + (1-T^-(1)-T^+(1))\,\rho(1,C)$

29-5

$\rho(1,C) = T^-(1)\rho(0,C) + T^+(1)\rho(2,C) + (1-T^-(1)-T^+(1))\,\rho(1,C)$

Solving recursion gives $\rho(1,C) = \left(\sum_{k=0}^{N-1}\prod_{i=1}^{k}\gamma^i\right)^{-1},$

$\gamma^i = T^-(i)/T^+(i)$

$T+(i)$= the probability that the i C strategists increases by 1

$T-(i)$ = the probability that the i C strategists decreases by 1

These two probabilities depend on the success (payoff) of each action

30

In python

```
def fixation(invader, resident):
    '''
    function for calculating the fixation
    probability of the invader
    in a population of residents
    '''
    result=0.
    for i in range(1,N):
        sub=1.
        for j in range(1,i+1):
            tmp=probIncreaseDecrease(j,invader, resident)
            sub*=(tmp[1]/float(tmp[0]))
        result += sub
    return np.clip(1./(1. + result), 0., 1.)
```

31

In python

```
def probIncreaseDecrease(k, invader, resident):
    '''
    This function calculates for a give number of invaders
     the probability
    that the number increases or decreases with one.
    '''
    fitvalue=fitness(k, invader, resident)
    increase=np.clip(((N-k)/float(N))*(k/float(N))*fermifunc(-beta,
                 fitvalue[1],fitvalue[0]),0.,1.)
    decrease=np.clip(((N-k)/float(N))*(k/float(N))*fermifunc(beta,
                 fitvalue[1],fitvalue[0]), 0., 1.)
    return [increase,decrease]
```

C Is the invader, k C and N-k D

$T+(k)$

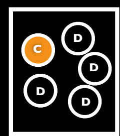

$T-(k)$

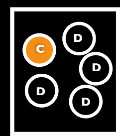

32-1

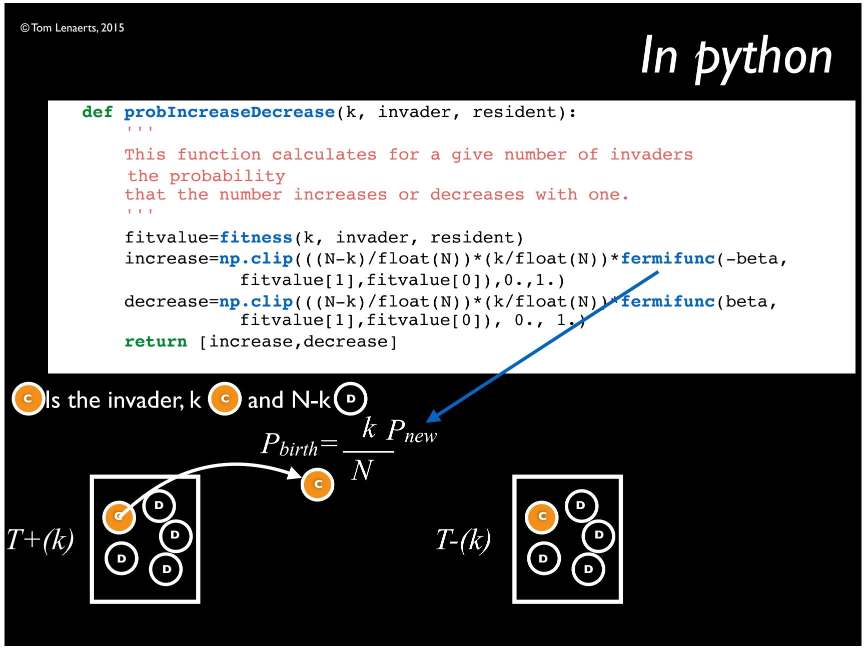

32-2

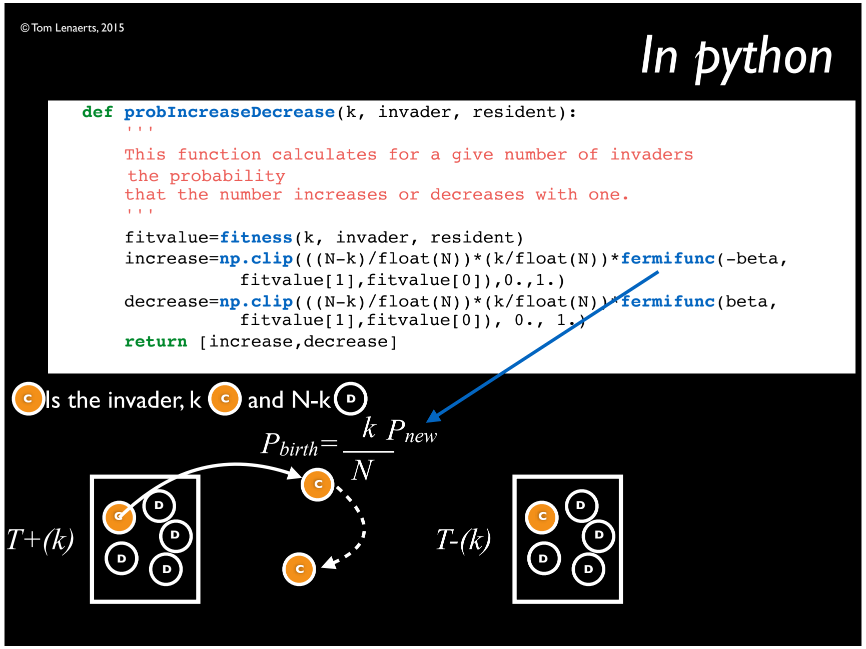

32-3

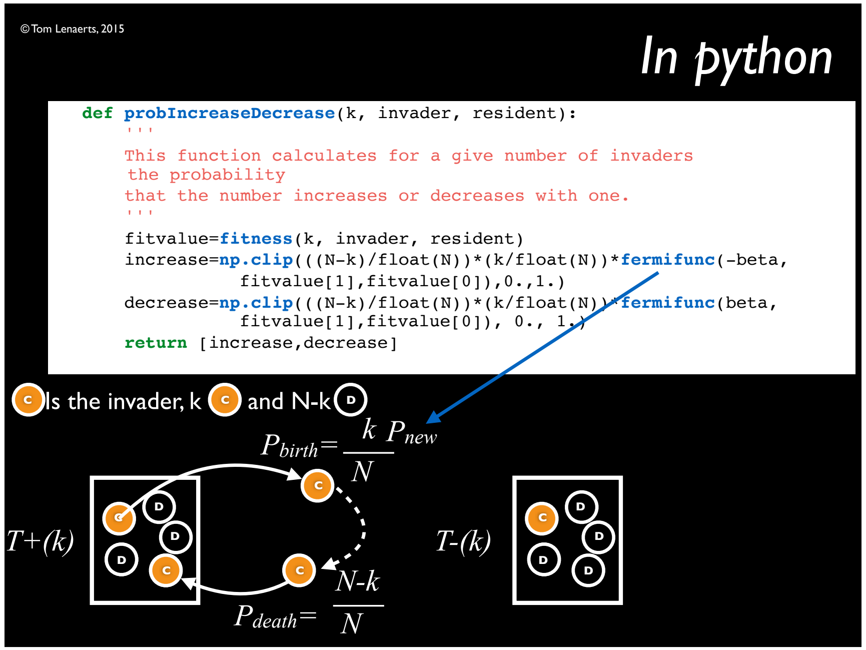

32-4

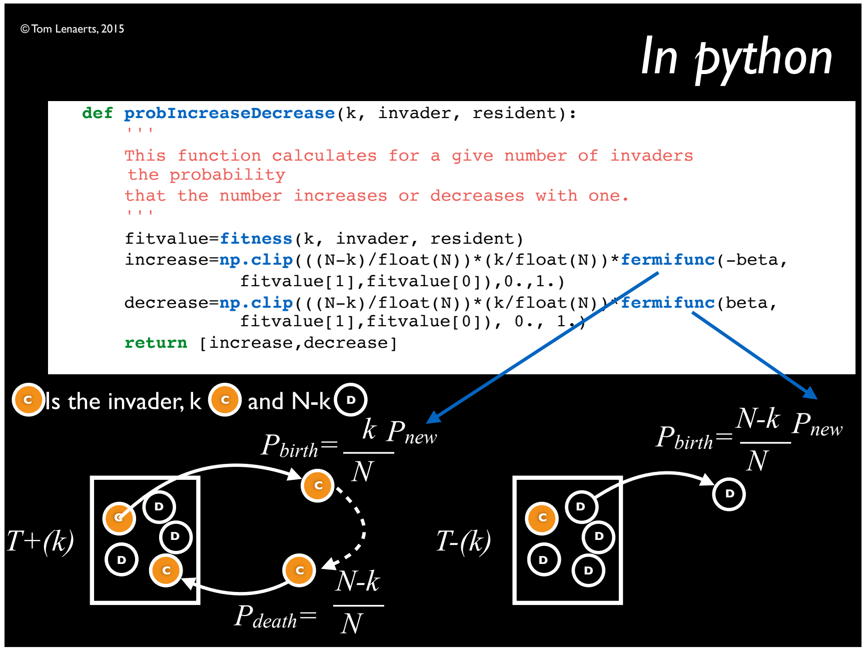

32-5

In python

```
def probIncreaseDecrease(k, invader, resident):
    '''
    This function calculates for a give number of invaders
     the probability
    that the number increases or decreases with one.
    '''
    fitvalue=fitness(k, invader, resident)
    increase=np.clip(((N-k)/float(N))*(k/float(N))*fermifunc(-beta,
             fitvalue[1],fitvalue[0]),0.,1.)
    decrease=np.clip(((N-k)/float(N))*(k/float(N))*fermifunc(beta,
             fitvalue[1],fitvalue[0]), 0., 1.)
    return [increase,decrease]
```

C Is the invader, k C and N-k D

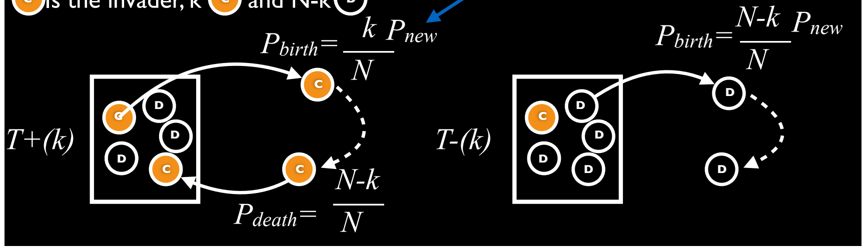

32-6

In python

```
def probIncreaseDecrease(k, invader, resident):
    '''
    This function calculates for a give number of invaders
     the probability
    that the number increases or decreases with one.
    '''
    fitvalue=fitness(k, invader, resident)
    increase=np.clip(((N-k)/float(N))*(k/float(N))*fermifunc(-beta,
             fitvalue[1],fitvalue[0]),0.,1.)
    decrease=np.clip(((N-k)/float(N))*(k/float(N))*fermifunc(beta,
             fitvalue[1],fitvalue[0]), 0., 1.)
    return [increase,decrease]
```

C Is the invader, k C and N-k D

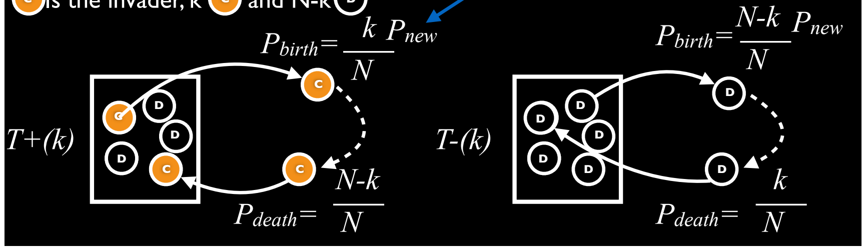

32-7

In python

```
def probIncreaseDecrease(k, invader, resident):
    '''
    This function calculates for a give number of invaders
     the probability
    that the number increases or decreases with one.
    '''
    fitvalue=fitness(k, invader, resident)
    increase=np.clip(((N-k)/float(N))*(k/float(N))*fermifunc(-beta,
             fitvalue[1],fitvalue[0]),0.,1.)
    decrease=np.clip(((N-k)/float(N))*(k/float(N))*fermifunc(beta,
             fitvalue[1],fitvalue[0]), 0., 1.)
    return [increase,decrease]
```

Remember Fermi $p=[1+e^{\beta(f(\mathbf{D})-f(\mathbf{C}))}]^{-1}$

```
def fermifunc(b,first, second):
    '''
    The probability that the first type imitates the second
    '''
    return 1./(1. + np.exp(-b*(first-second)))
```

33

In python

```
def fitness(k, invader, resident):
    '''
    The fitness function determines the average payoff of k
     invaders and N-k residents
    in the popiulation of N players.
    '''
    resultA=(((k-1)*payoff[invader][invader])+
       ((N-k)*payoff[invader, resident]))/float(N-1)
    resultB=((k*payoff[resident][invader])+
       ((N-k-1)*payoff[resident, resident]))/float(N-1)
    return [resultA, resultB]
```

resultA = average payoff of an invader playing against his own type A or the other type B

resultB = average payoff of an invader playing against his own type B or the other type A

34

©Tom Lenaerts, 2015

In python

```
def transition_and_fixprob_matrix():
    transitions=np.zeros((q,q))
    fixprobs=np.zeros((q,q))
    for first in range(q):
        transitions[first,first]=1.
        for second in range(q) :
            if second != first :
                fp=fixation(second, first)
                fixprobs[first][second]=(fp/drift)
                transitions[first][second]=fp/float(q-1)
                transitions[first][first]=
                 transitions[first][first]-(fp/float(q-1))
    return [transitions,fixprobs]
```

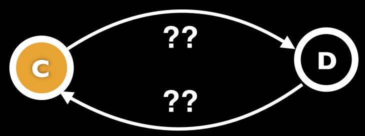

35

©Tom Lenaerts, 2015

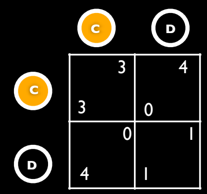

with β=0.01

ρ_N=drift=1/N

```
N=100 #Population size
T=3. #Temptation to defect
R=4. #Reward for mutual cooperation
P=0. #Punishment for mutual defection
S=1. #Suckers payoff for unilateral cooperation
q=2 #Number of strategies
drift=1.0/N

strats=['C','D']
payoff=np.array([R,S,T,P]).reshape(2,2)
beta=0.01

t,f=transition_and_fixprob_matrix()
print ("transition probabilities (first %s,
second %s and read as row invading column) \n
%s" %(strats[0], strats[1],f))
```

```
transition probabilities (first
C, second D and read as row
invading column)
 [[ 0.          0.5951908 ]
 [ 1.55445772  0.         ]]
```

36

©Tom Lenaerts, 2015

Stationary distribution

Which produces a **reduced Markov chain**

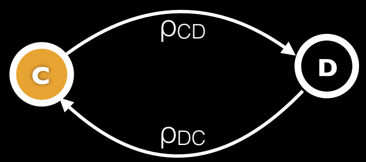

Fudenberg, D., & Imhof, L.A. (2006). Imitation processes with small mutations. J. Econ.Theo, 131,251–262.

Imhof, L.A., Fudenberg, D., & Nowak, M.A. (2005). Evolutionary cycles of cooperation and defection. Proc Nat Acad Sci USA, 102(31), 10797–10800.

37-1

©Tom Lenaerts, 2015

Stationary distribution

Which produces a **reduced Markov chain**

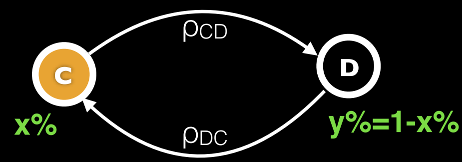

For which the **stationary distributions** can be calculated = *how likely it is to end up in either monomorphic state*

Fudenberg, D., & Imhof, L.A. (2006). Imitation processes with small mutations. J. Econ.Theo, 131,251–262.

Imhof, L.A., Fudenberg, D., & Nowak, M.A. (2005). Evolutionary cycles of cooperation and defection. Proc Nat Acad Sci USA, 102(31), 10797–10800.

37-2

```python
#calculate stationary distributions using eigenvalues and eigenvectors
w,v=np.linalg.eig(t.transpose())

#look for the element closest to 1 in the list of eigenvalues
j_stationary=np.argmin(abs(w-1.0))

# the, is essential to access the matrix by column
p_stationary=abs(v[:,j_stationary].real)

#normalize
p_stationary /= p_stationary.sum()

print ("stationary distribution (first %s, second %s) \n %s" %(strats[0],
strats[1], p_stationary))
```

The stationary distribution is the **minimum left eigenvalue** of the transition matrix.

38

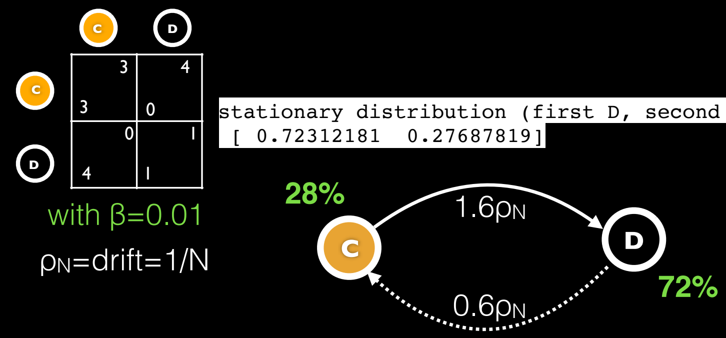


```
stationary distribution (first D, second C)
 [ 0.72312181  0.27687819]
```

with β=0.01

ρ_N=drift=1/N

Python **code for numerical approximation** (random walks in Markov chain) see 2IPD-numerical.ipynb (see assignment 1)

39

For varying β

```python
#calculate stationary distributions using eigenvalues and eigenvectors
distribution=[]
betas=[0.00001,0.0001, 0.001, 0.01, 0.1, 1.0, 10.0, 100.0]
for beta in betas:
    t,f=transition_and_fixprob_matrix()
    w,v=np.linalg.eig(t.transpose())
    j_stationary=np.argmin(abs(w-1.0))
    p_stationary=abs(v[:,j_stationary].real)
    p_stationary /= p_stationary.sum()
    distribution.append(p_stationary)

darray=np.asarray(distribution)

# and plot the curve
```

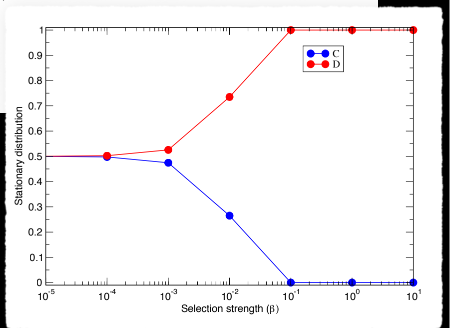

40-1

©Tom Lenaerts, 2015

For varying β

```python
#calculate stationary distributions using eigenvalues and eigenvectors
distribution=[]
betas=[0.00001,0.0001, 0.001, 0.01, 0.1, 1.0, 10.0, 100.0]
for beta in betas:
    t,f=transition_and_fixprob_matrix()
    w,v=np.linalg.eig(t.transpose())
    j_stationary=np.argmin(abs(w-1.0))
    p_stationary=abs(v[:,j_stationary].real)
    p_stationary /= p_stationary.sum()
    distribution.append(p_stationary)

darray=np.asarray(distribution)

# and plot the curve
```

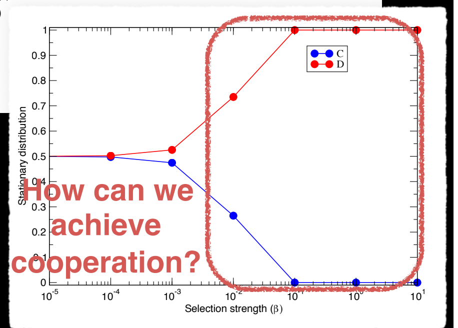

40-2

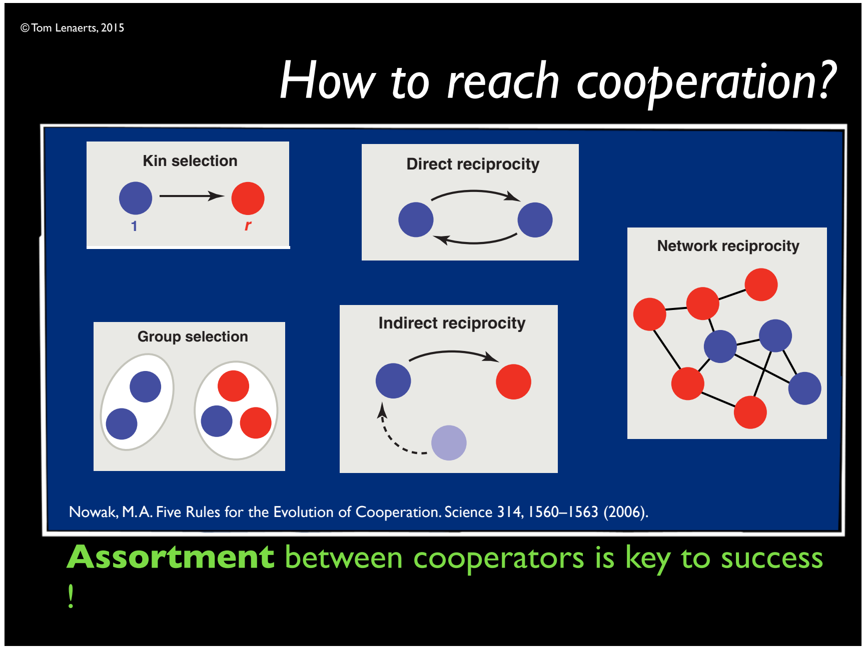
©Tom Lenaerts, 2015
How to reach cooperation?
Kin selection
1
r
Direct reciprocity
Network reciprocity
Group selection
Indirect reciprocity
Nowak, M.A. Five Rules for the Evolution of Cooperation. Science 314, 1560–1563 (2006).
Assortment between cooperators is key to success !

41

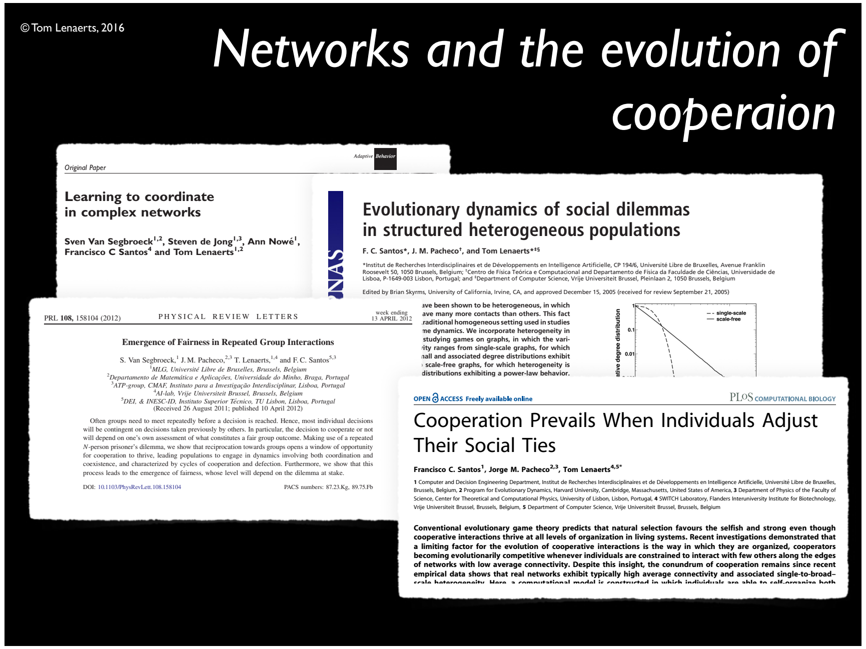
©Tom Lenaerts, 2016
Networks and the evolution of cooperaion
Learning to coordinate in complex networks
Evolutionary dynamics of social dilemmas in structured heterogeneous populations
PHYSICAL REVIEW LETTERS
Emergence of Fairness in Repeated Group Interactions
Cooperation Prevails When Individuals Adjust Their Social Ties

42

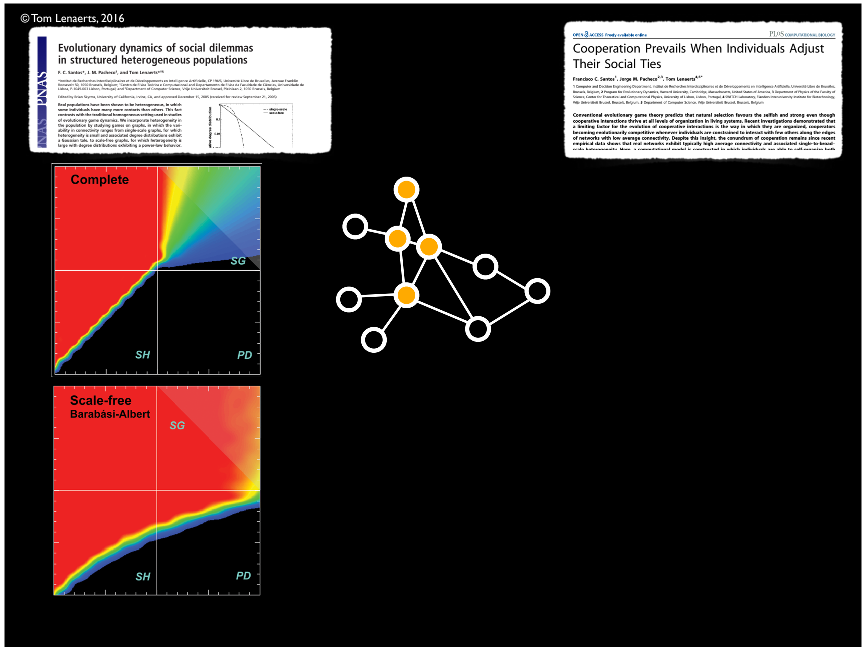
©Tom Lenaerts, 2016
Evolutionary dynamics of social dilemmas in structured heterogeneous populations
Cooperation Prevails When Individuals Adjust Their Social Ties
Complete
SG
SH
PD
Scale-free
Barabási-Albert
SG
SH
PD

43-1

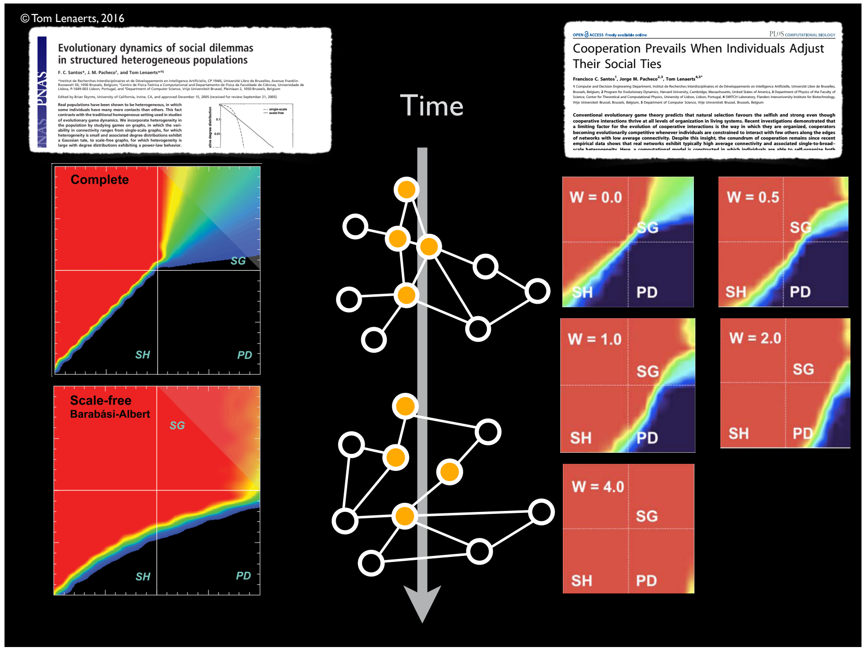
©Tom Lenaerts, 2016
Evolutionary dynamics of social dilemmas in structured heterogeneous populations
Cooperation Prevails When Individuals Adjust Their Social Ties
Time
Complete
SG
SH
PD
Scale-free
Barabási-Albert
SG
SH
PD
W = 0.0
W = 0.5
W = 1.0
W = 2.0
W = 4.0

43-2

Or more elaborate behaviours


Mechanisms of Social Cognition

Chris D. Frith[1,3] and Uta Frith[2,3]

44

Trust drives social interactions …

How to make sure we can trust a partner?

45

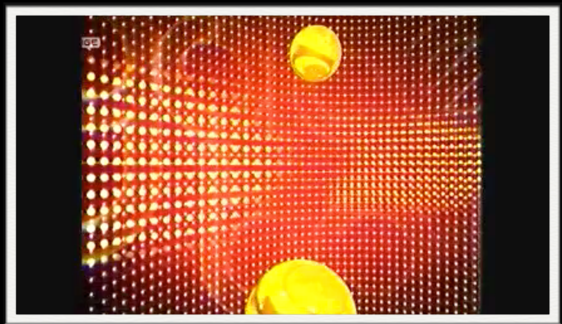

46-1

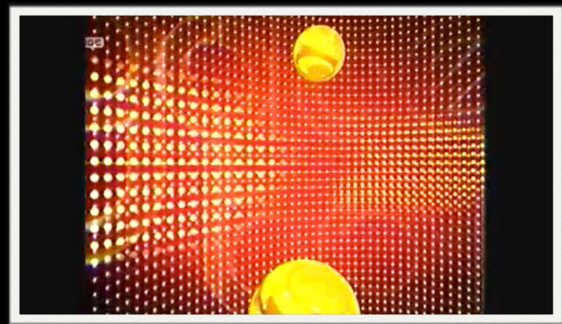

46-2

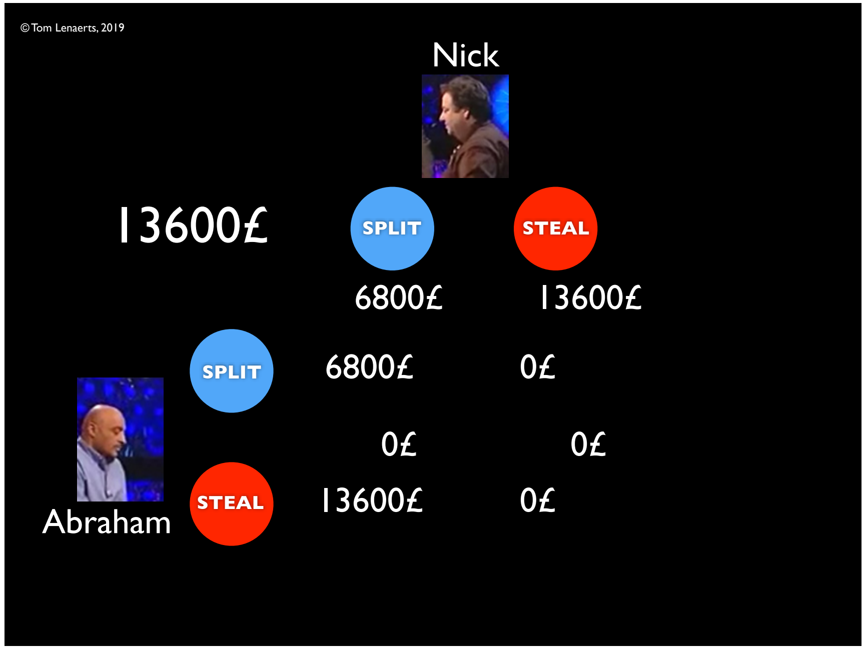

47-1

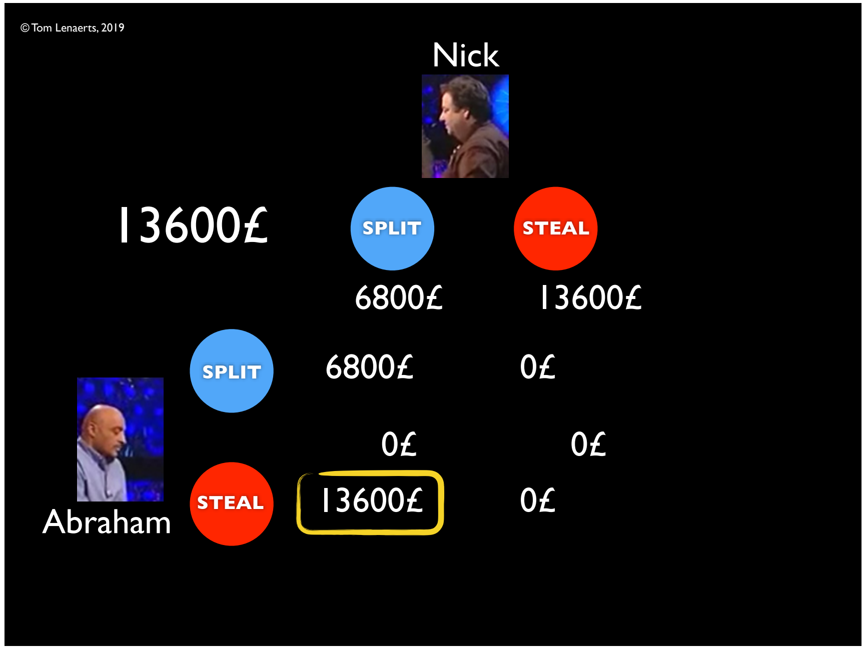

47-2

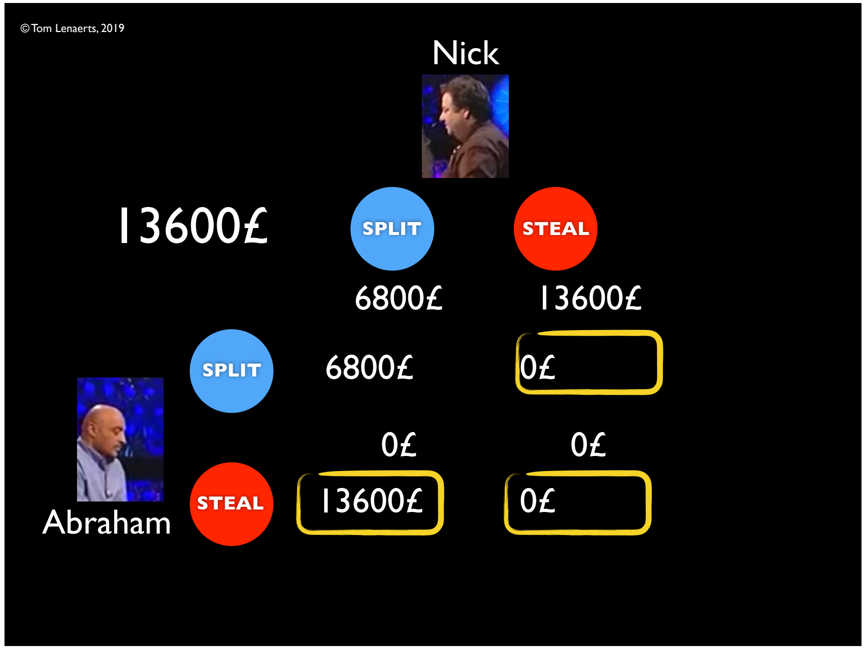

47-3

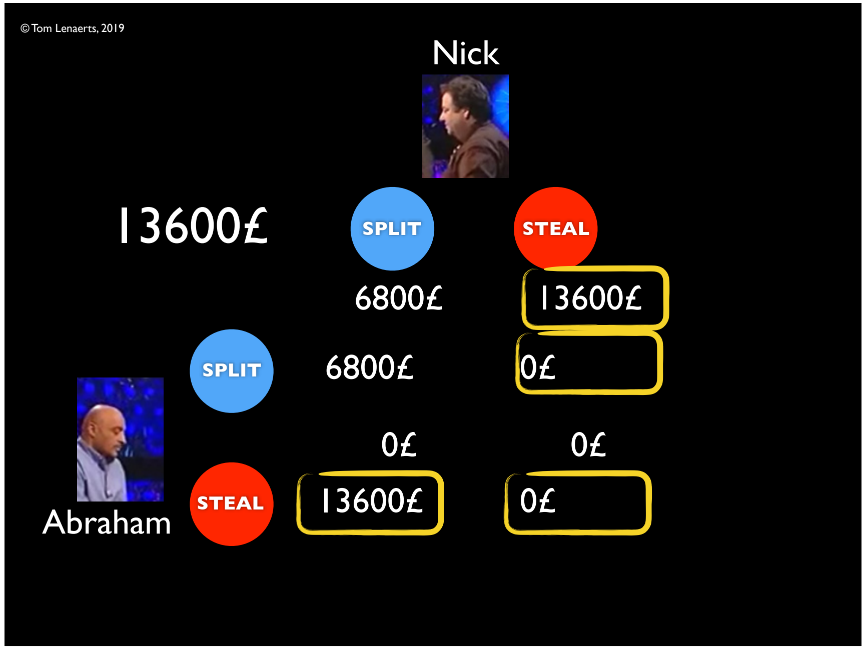

47-4

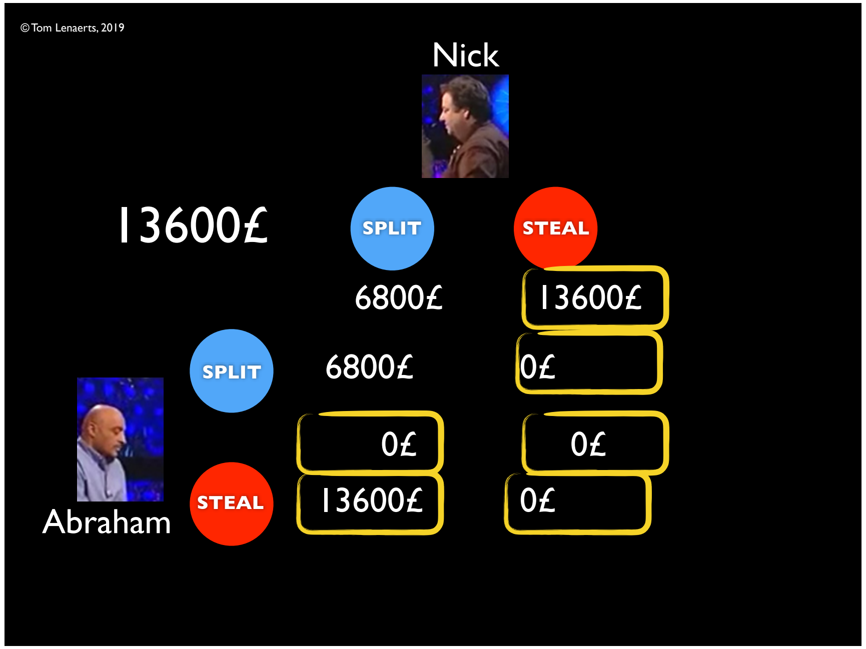

47-5

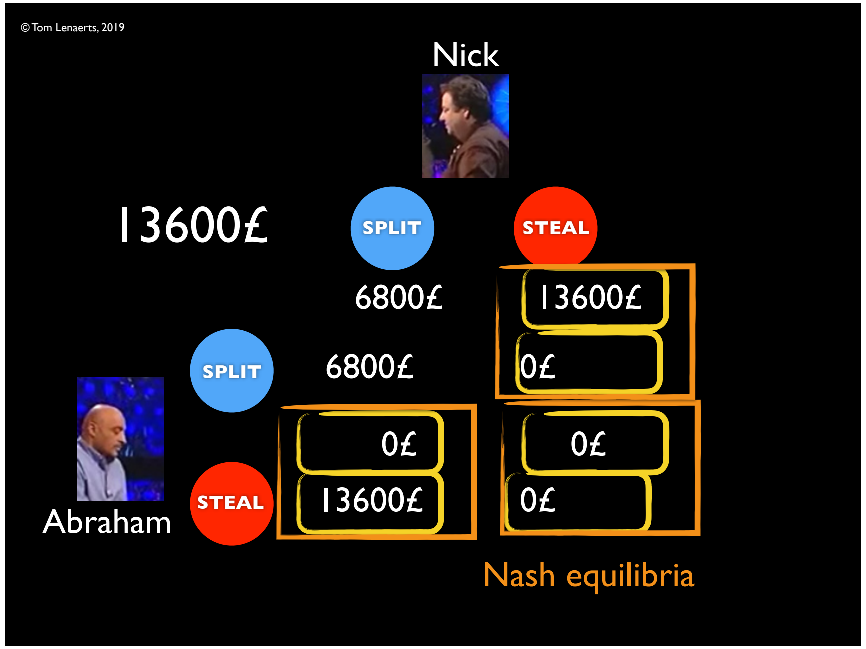

47-6

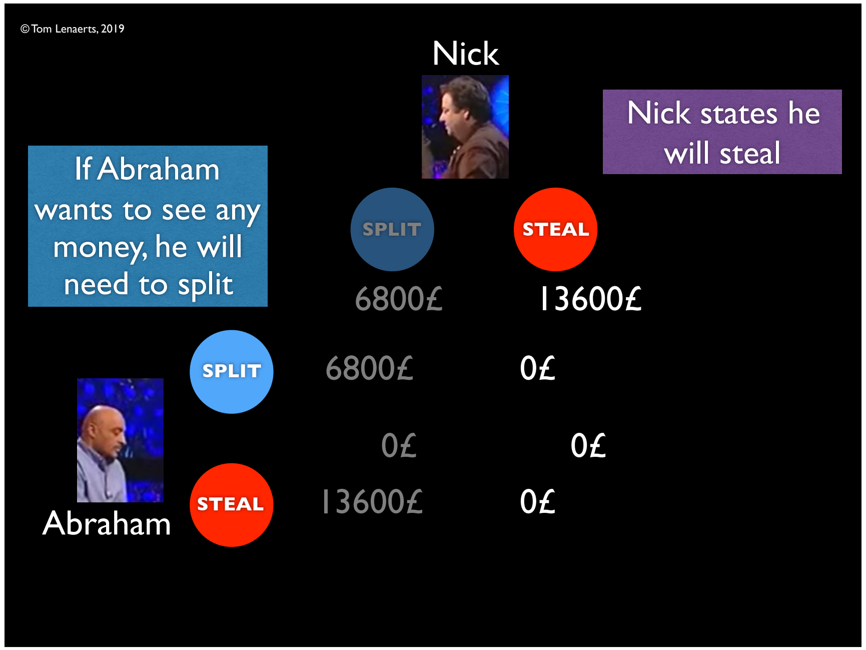

48-1

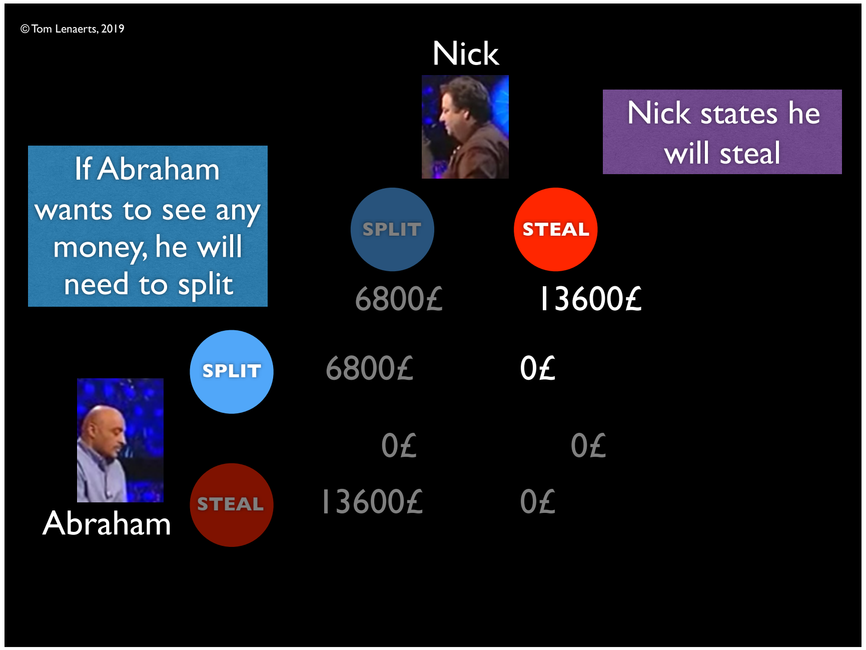

48-2

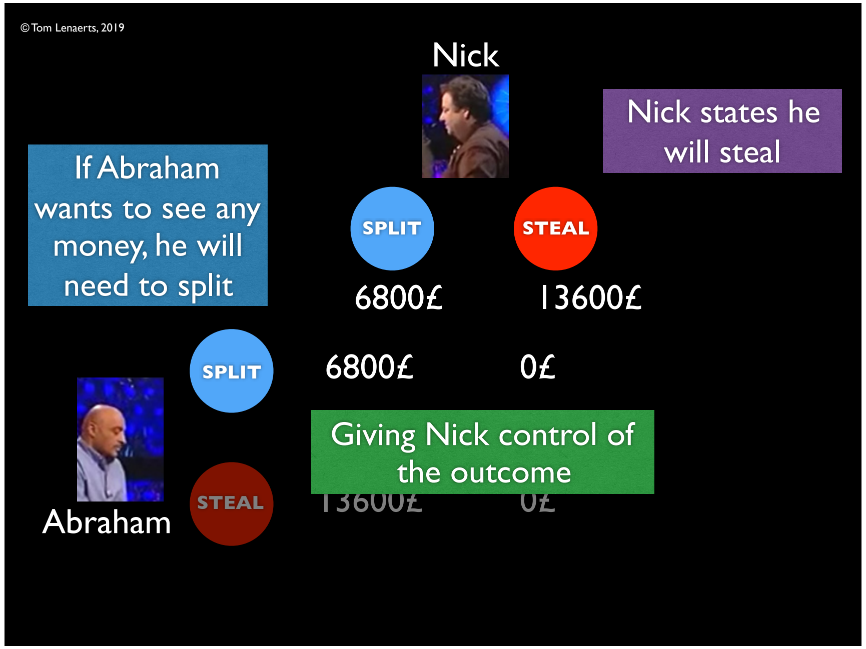

48-3

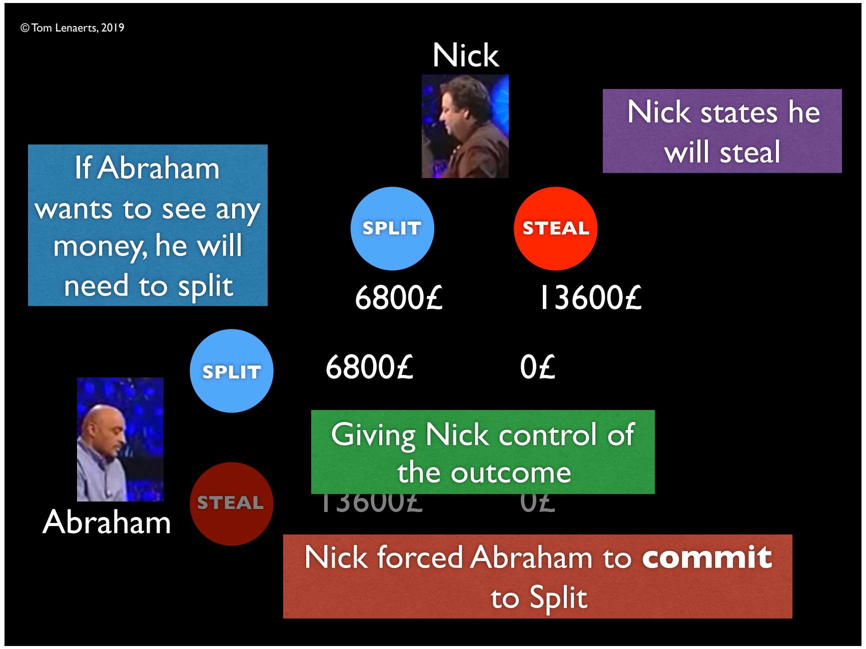

48-4

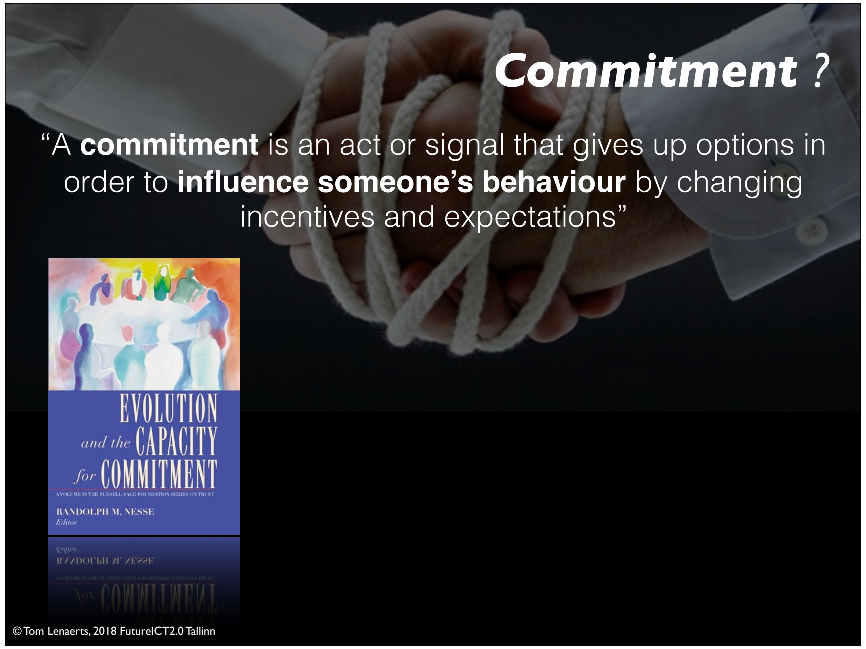

49-1

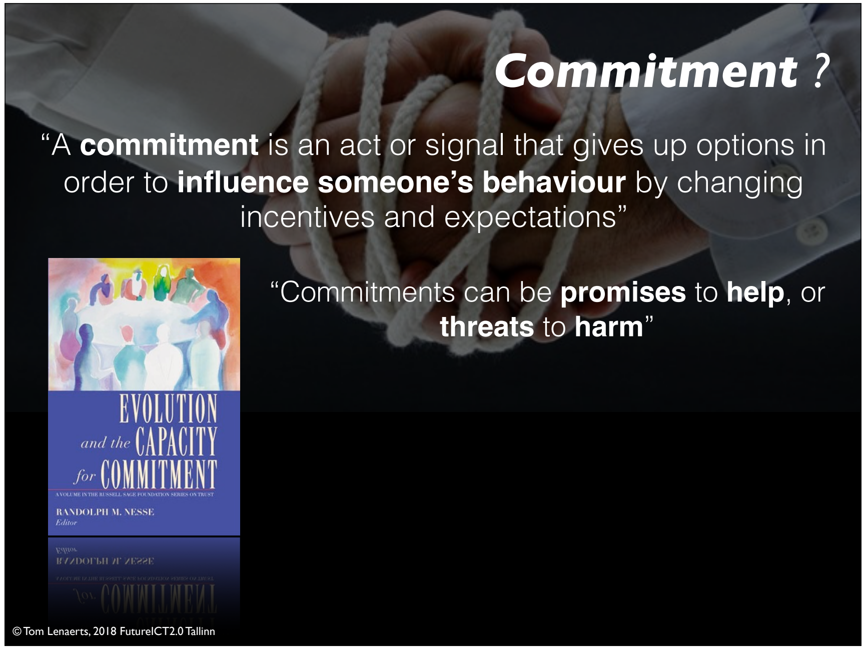

49-2

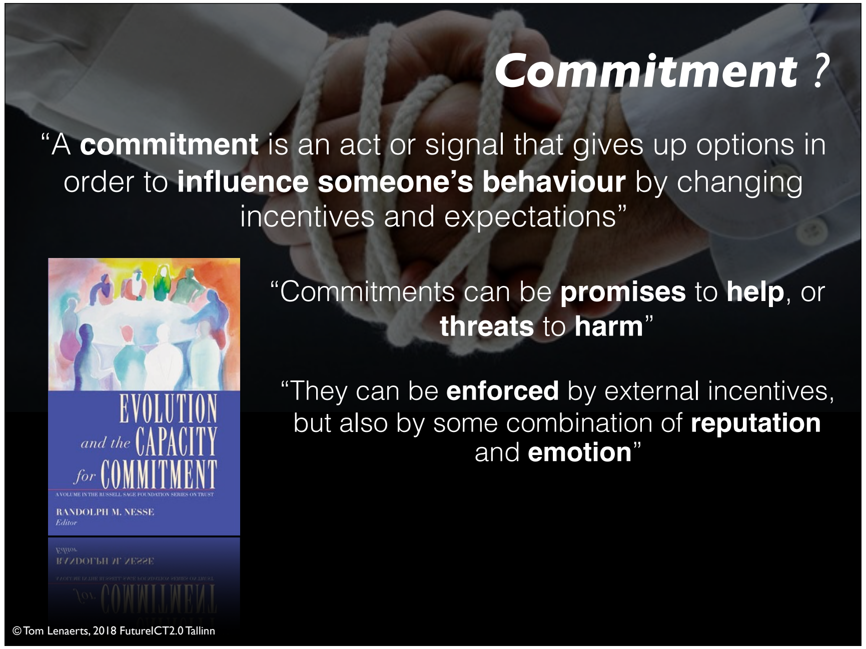

49-3

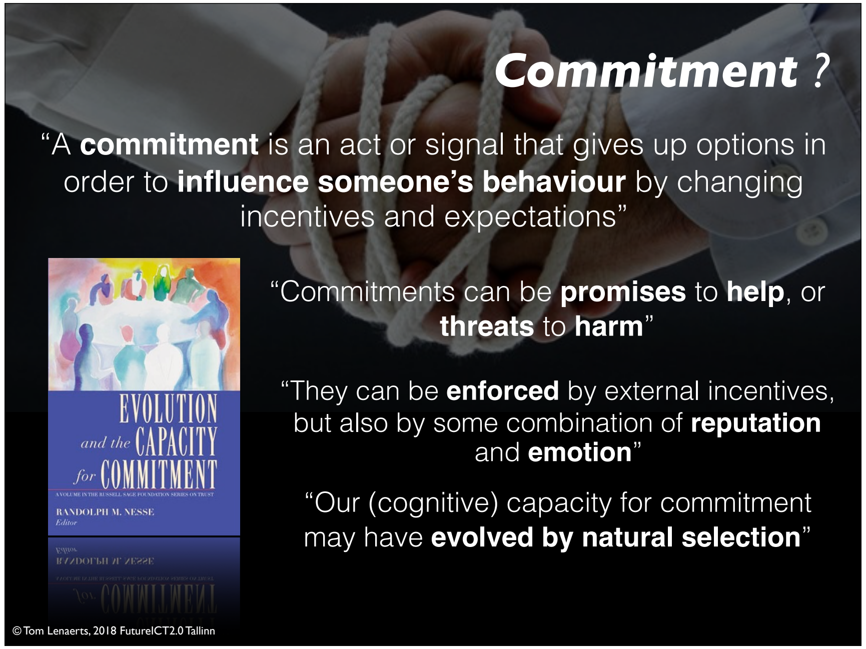

49-4

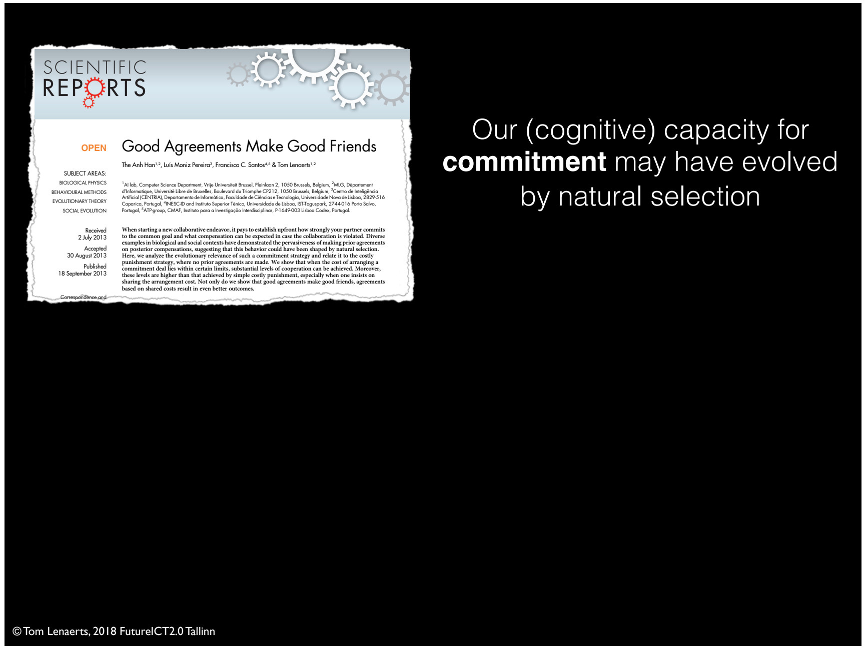

50-1

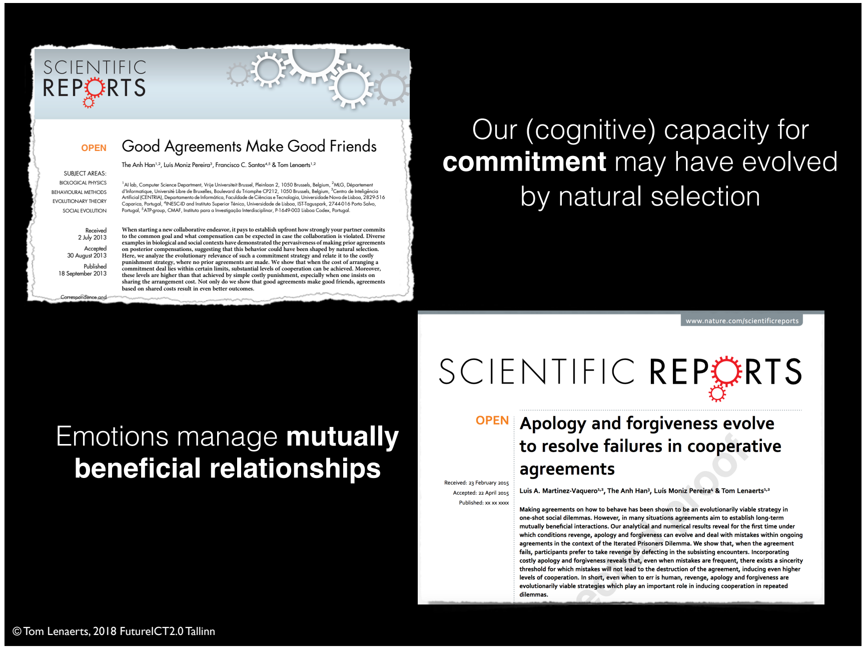

50-2

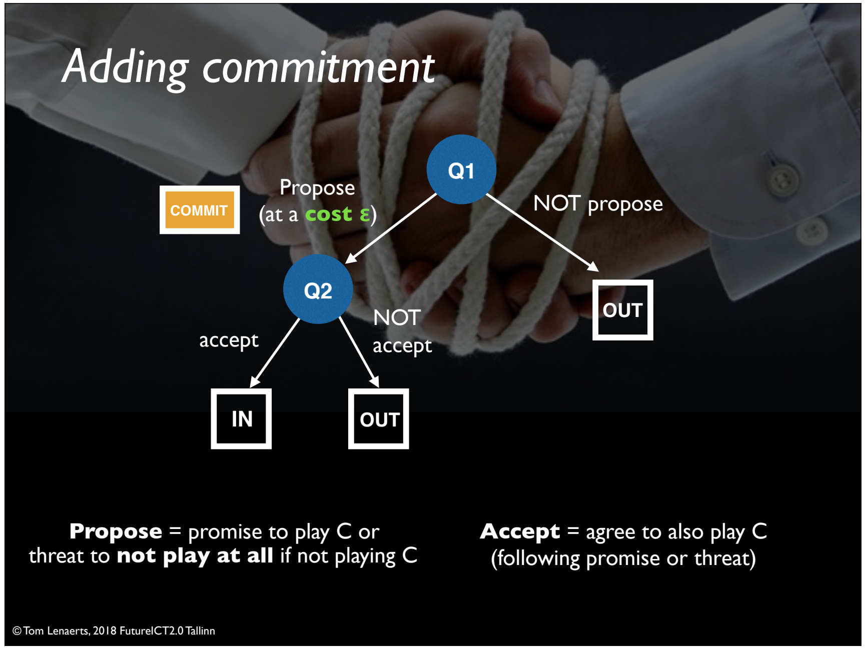

51-1

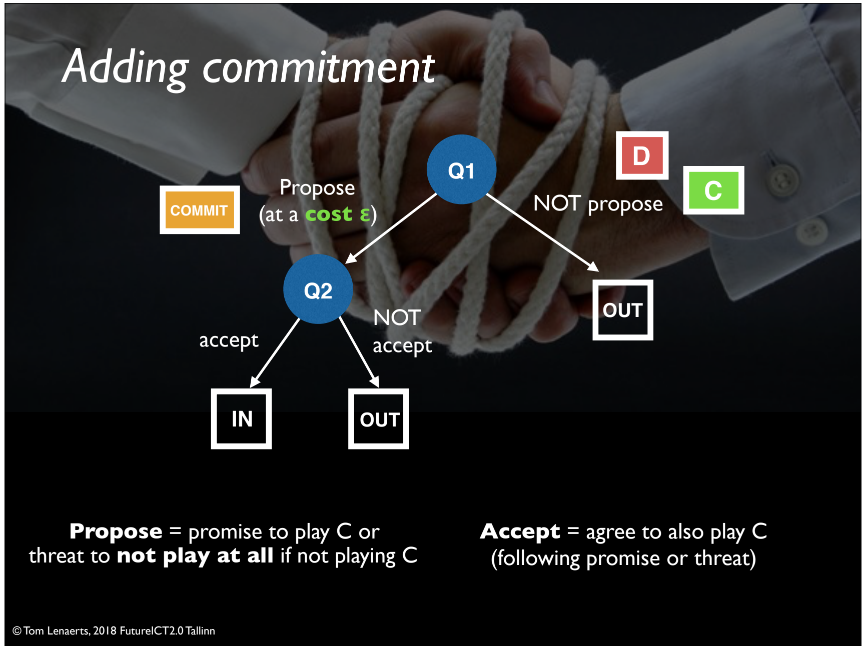

51-2

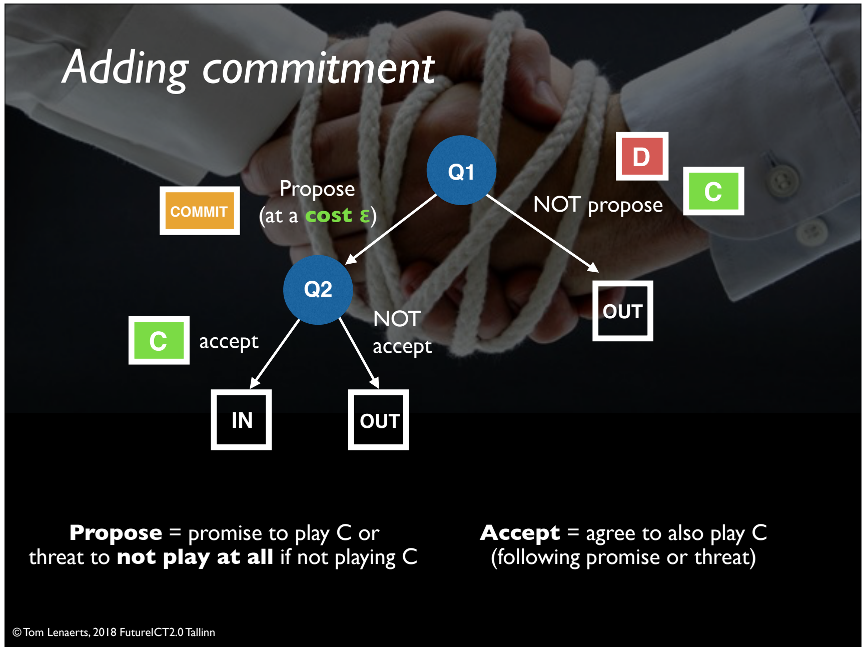

51-3

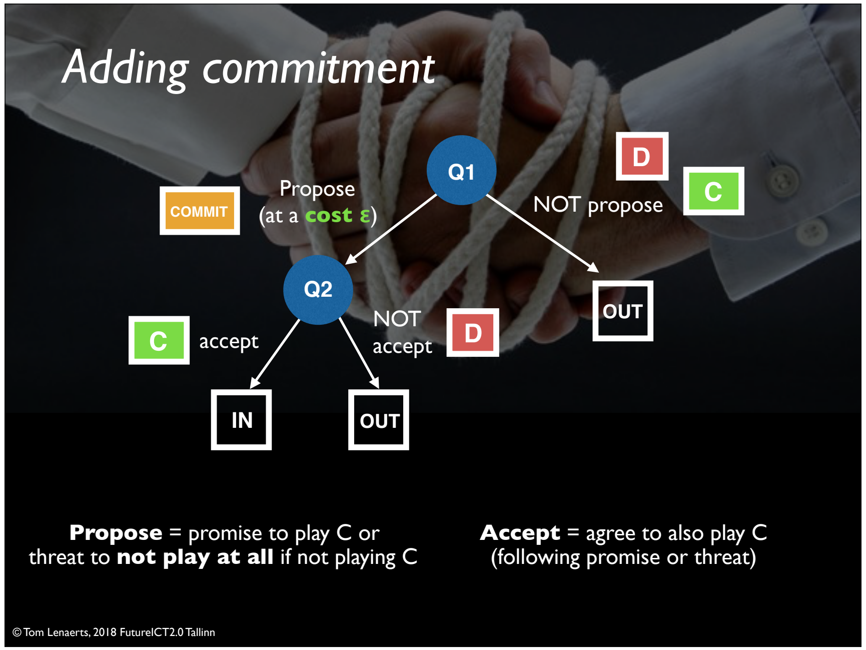

51-4

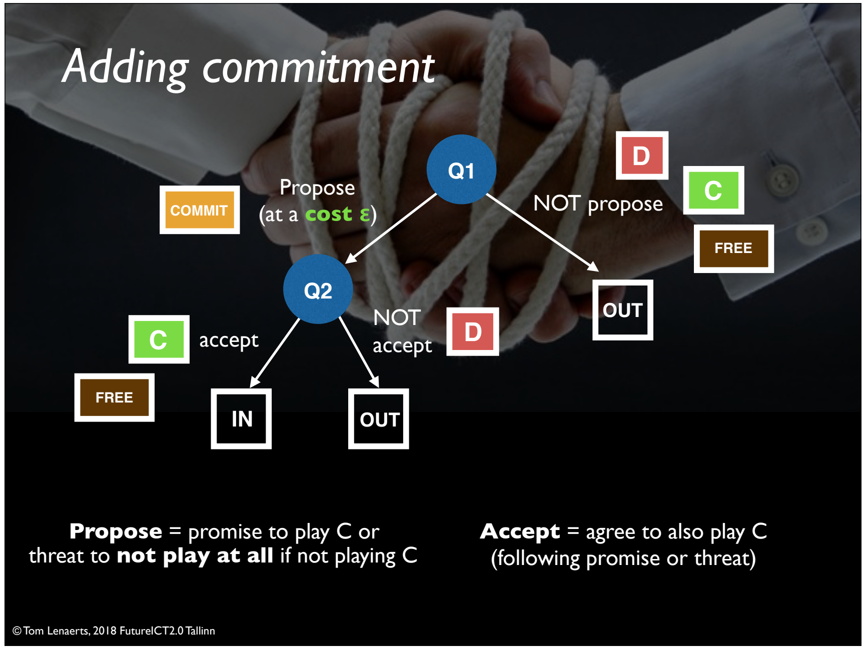

51-5

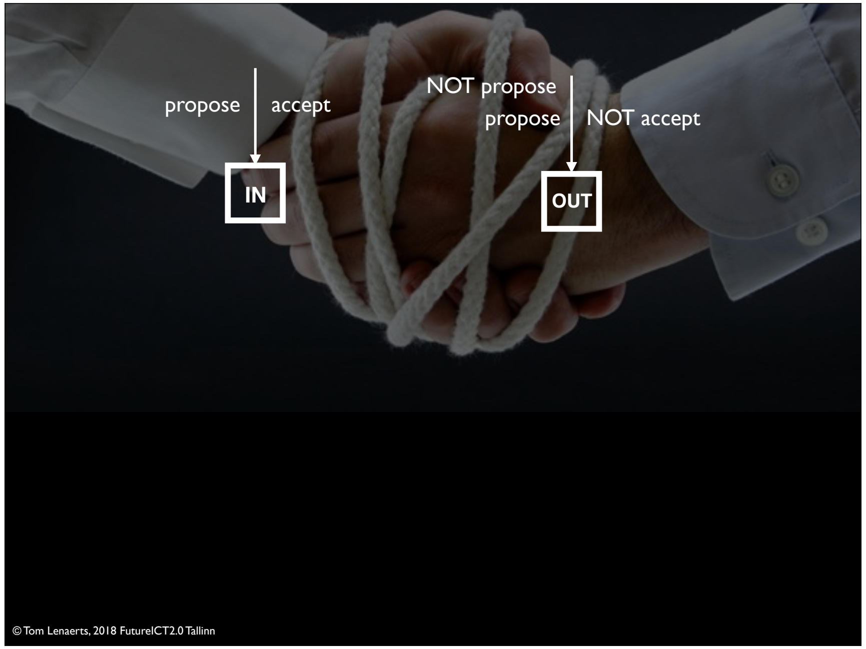

52-1

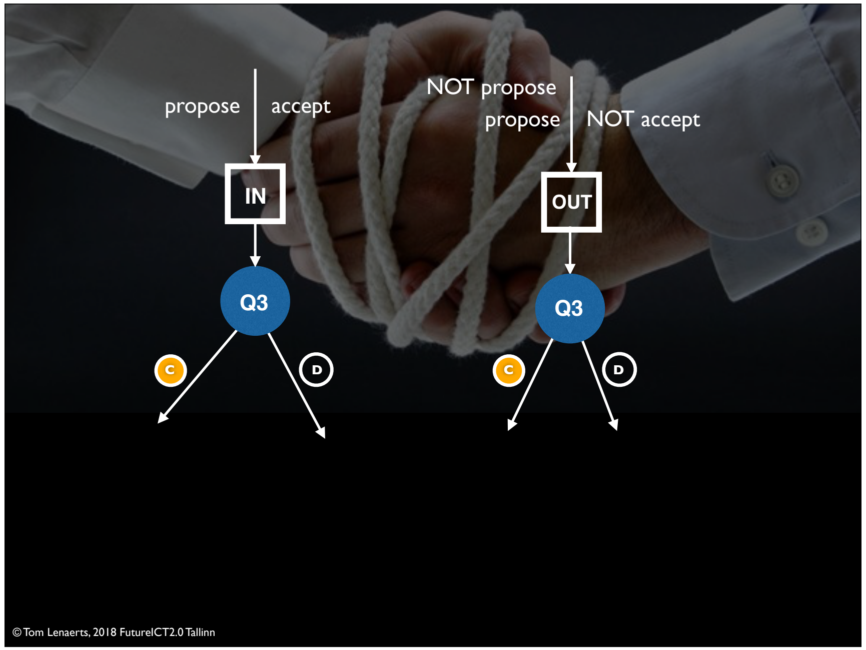

52-2

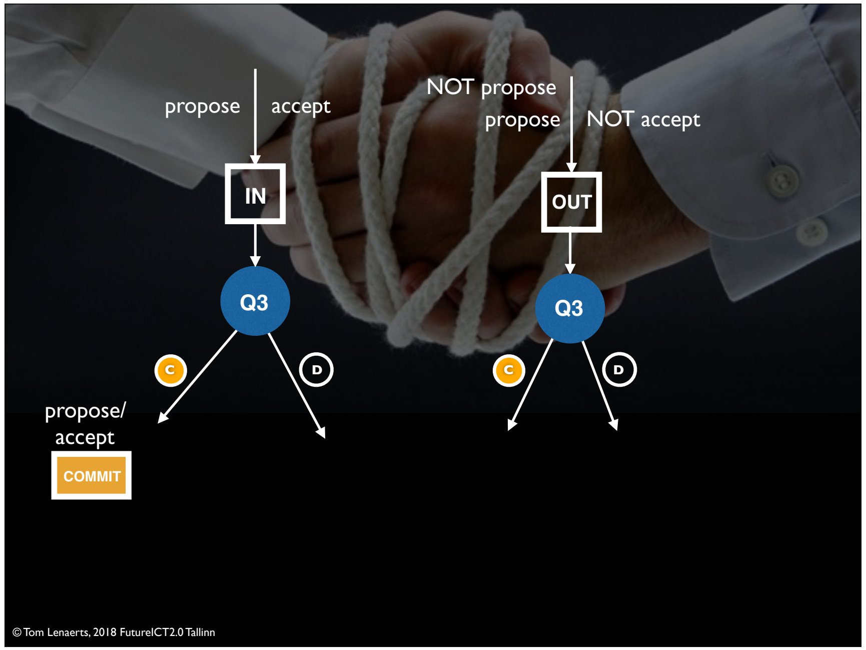

52-3

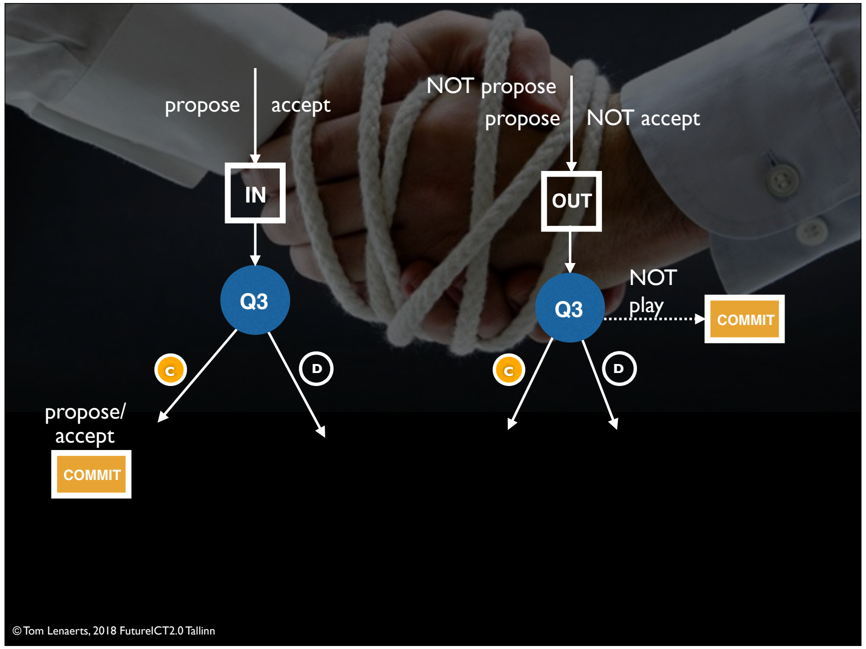

52-4

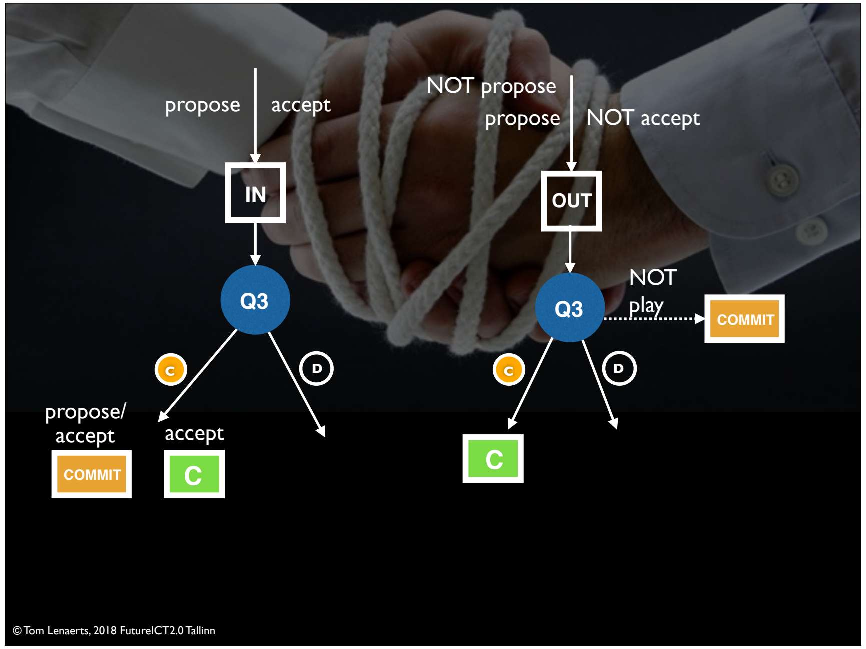

52-5

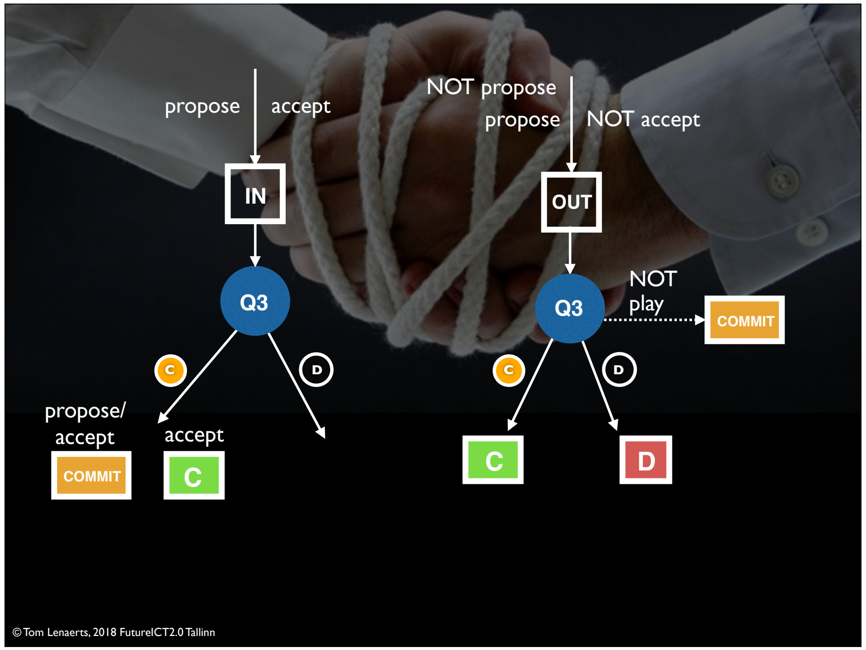

52-6

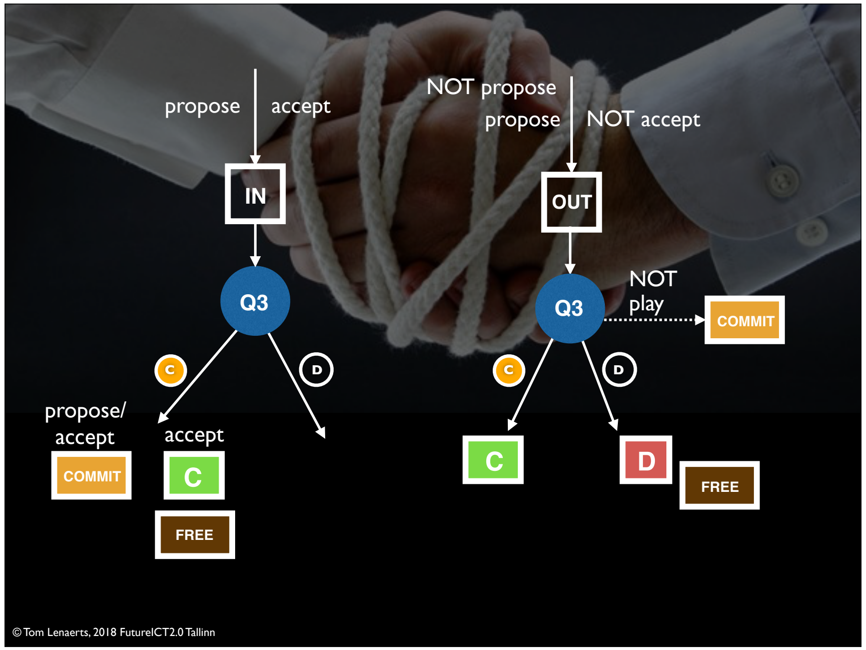

52-7

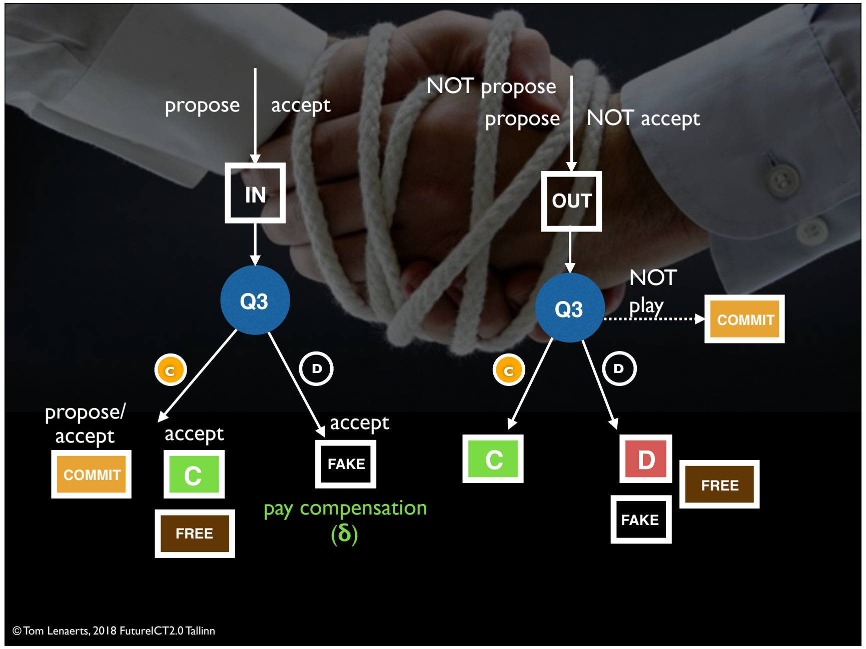

52-8

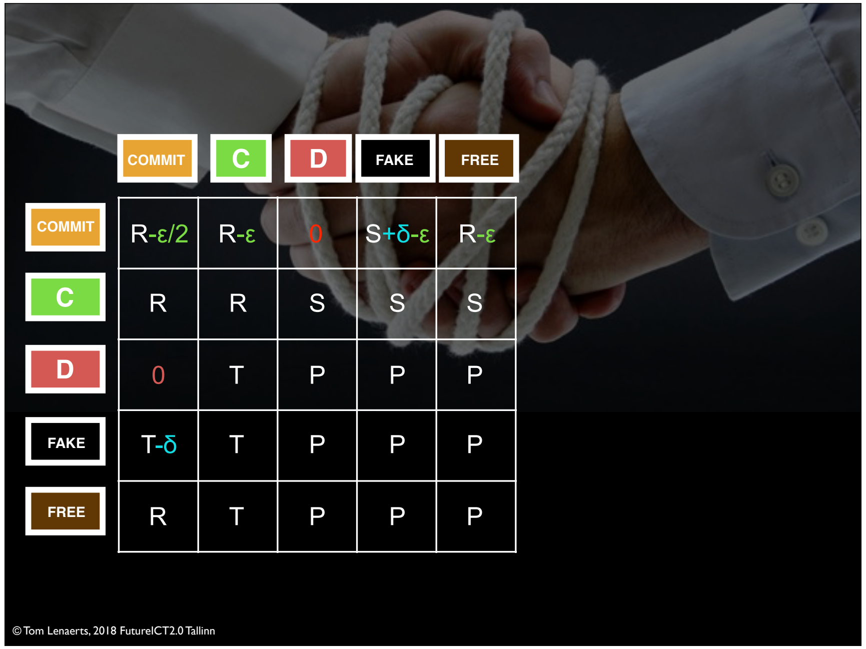

	COMMIT	C	D	FAKE	FREE
COMMIT	R-ε/2	R-ε	0	S+δ-ε	R-ε
C	R	R	S	S	S
D	0	T	P	P	P
FAKE	T-δ	T	P	P	P
FREE	R	T	P	P	P

53-1

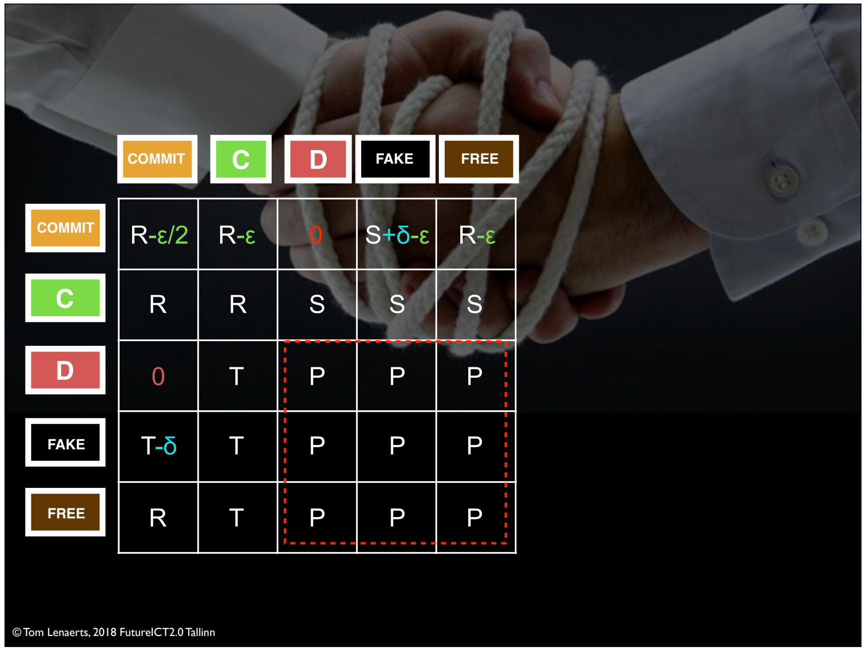

	COMMIT	C	D	FAKE	FREE
COMMIT	R-ε/2	R-ε	0	S+δ-ε	R-ε
C	R	R	S	S	S
D	0	T	P	P	P
FAKE	T-δ	T	P	P	P
FREE	R	T	P	P	P

53-2

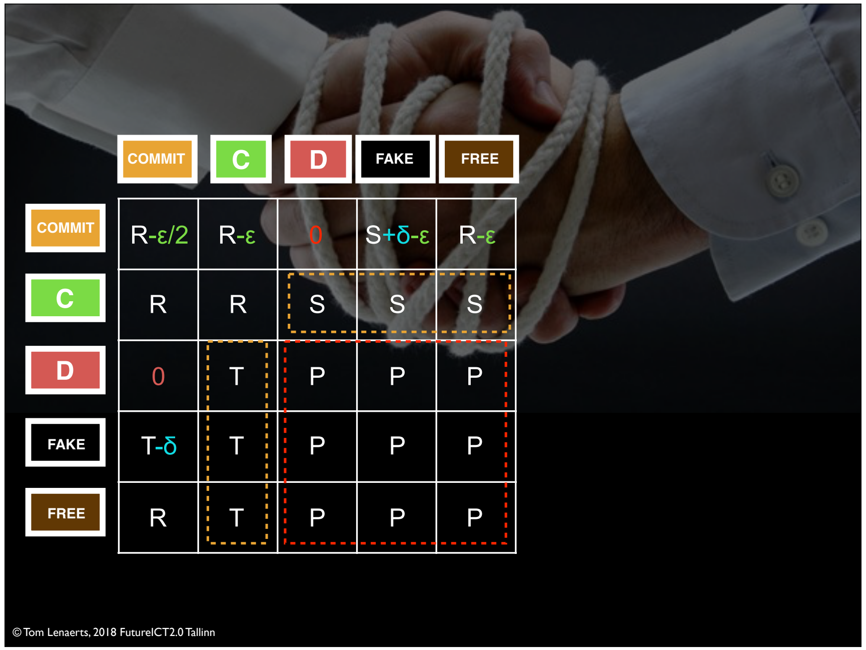

	COMMIT	C	D	FAKE	FREE
COMMIT	R-ε/2	R-ε	0	S+δ-ε	R-ε
C	R	R	S	S	S
D	0	T	P	P	P
FAKE	T-δ	T	P	P	P
FREE	R	T	P	P	P

53-3

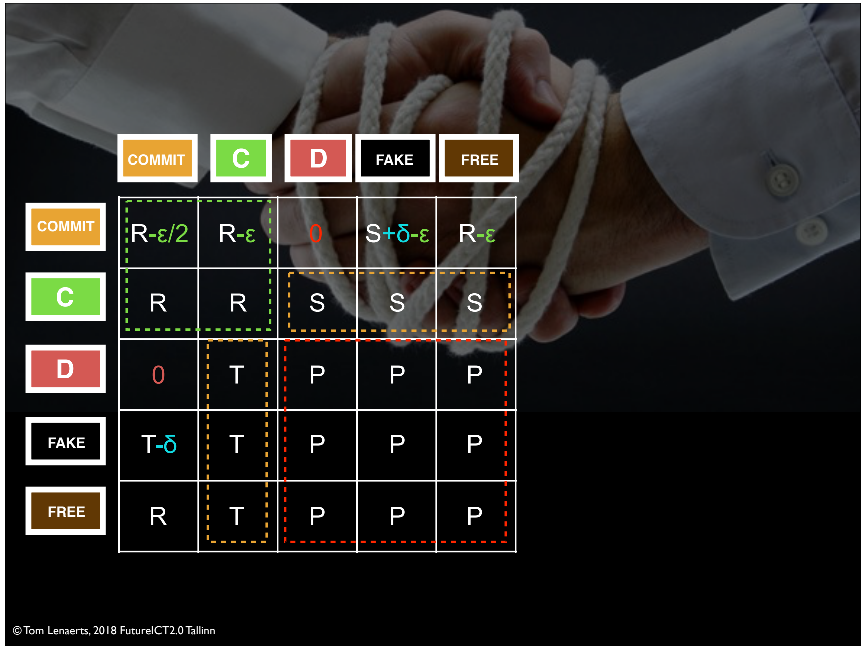

	COMMIT	C	D	FAKE	FREE
COMMIT	R-ε/2	R-ε	0	S+δ-ε	R-ε
C	R	R	S	S	S
D	0	T	P	P	P
FAKE	T-δ	T	P	P	P
FREE	R	T	P	P	P

53-4

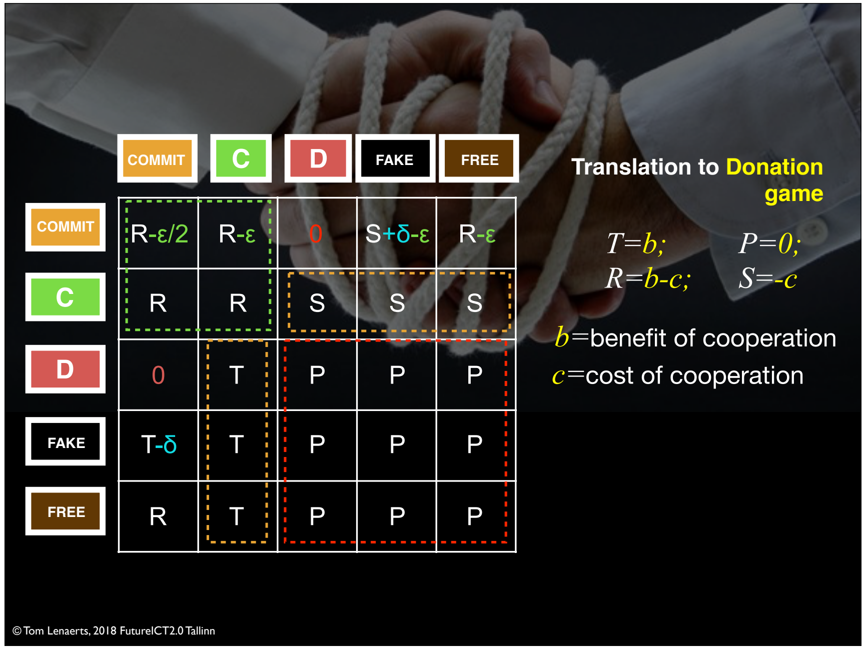

	COMMIT	C	D	FAKE	FREE
COMMIT	R-ε/2	R-ε	0	S+δ-ε	R-ε
C	R	R	S	S	S
D	0	T	P	P	P
FAKE	T-δ	T	P	P	P
FREE	R	T	P	P	P

53-5

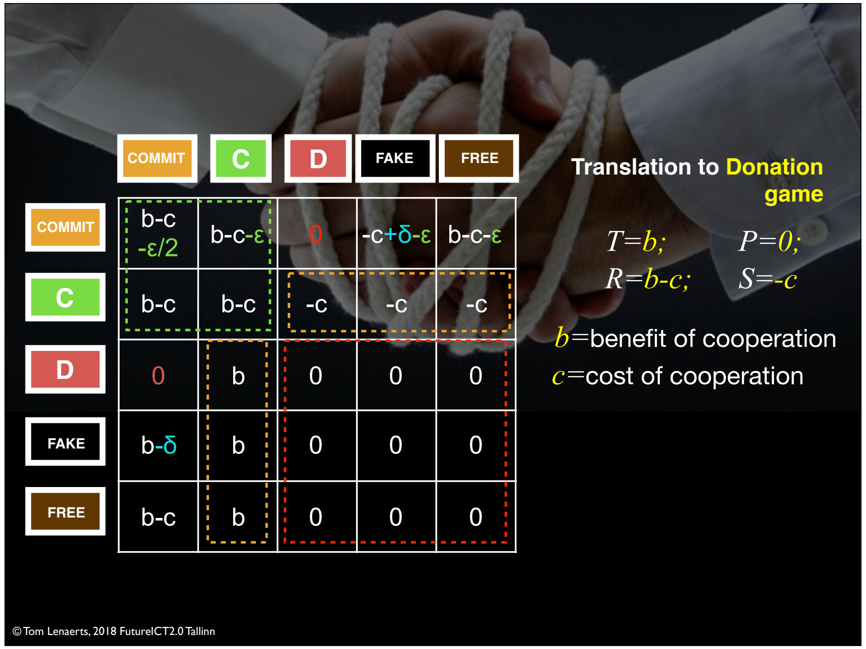

	COMMIT	C	D	FAKE	FREE
COMMIT	b-c -ε/2	b-c-ε	0	-c+δ-ε	b-c-ε
C	b-c	b-c	-c	-c	-c
D	0	b	0	0	0
FAKE	b-δ	b	0	0	0
FREE	b-c	b	0	0	0

54-1

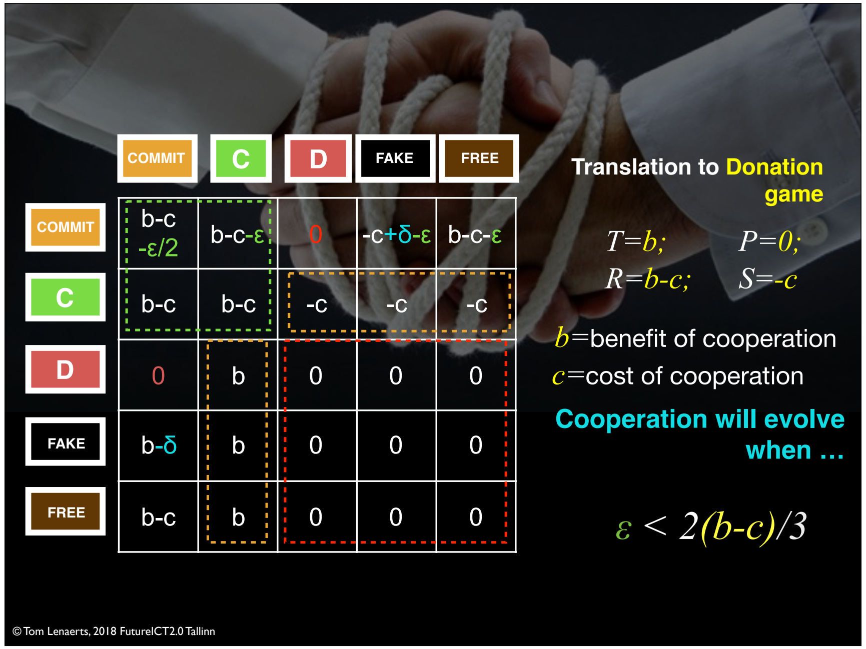

	COMMIT	C	D	FAKE	FREE
COMMIT	b-c -ε/2	b-c-ε	0	-c+δ-ε	b-c-ε
C	b-c	b-c	-c	-c	-c
D	0	b	0	0	0
FAKE	b-δ	b	0	0	0
FREE	b-c	b	0	0	0

54-2

Translation to Donation game

	COMMIT	C	D	FAKE	FREE
COMMIT	b-c -ε/2	b-c-ε	0	-c+δ-ε	b-c-ε
C	b-c	b-c	-c	-c	-c
D	0	b	0	0	0
FAKE	b-δ	b	0	0	0
FREE	b-c	b	0	0	0

$T=b$; $P=0$;
$R=b-c$; $S=-c$

b=benefit of cooperation
c=cost of cooperation

Cooperation will evolve when …

$\varepsilon < 2(b-c)/3$

$\delta > c+3\varepsilon/4$

54-3

See Jupyter notebook

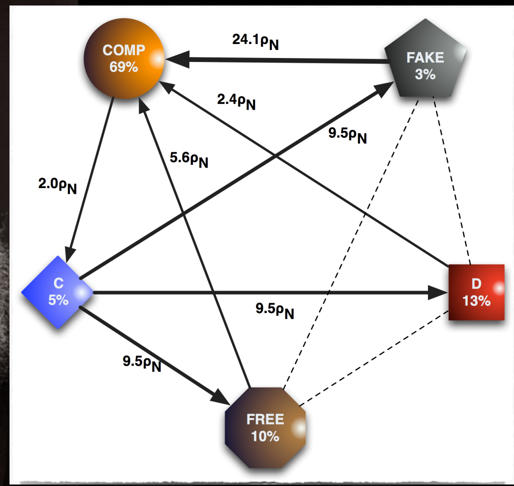

Same conclusions hold for 2 or more players

55

See Jupyter notebook

β=0.1

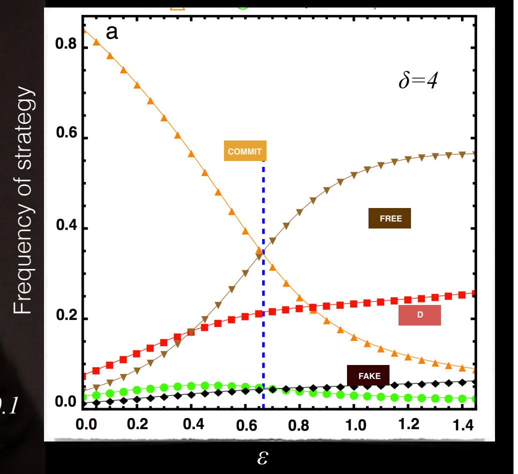

Same conclusions hold for 2 or more players

56

*Commitments are often **long-term***

With a probability ω the interaction continues

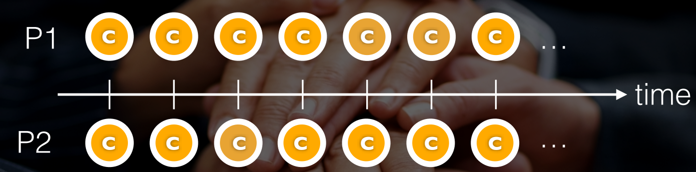

57

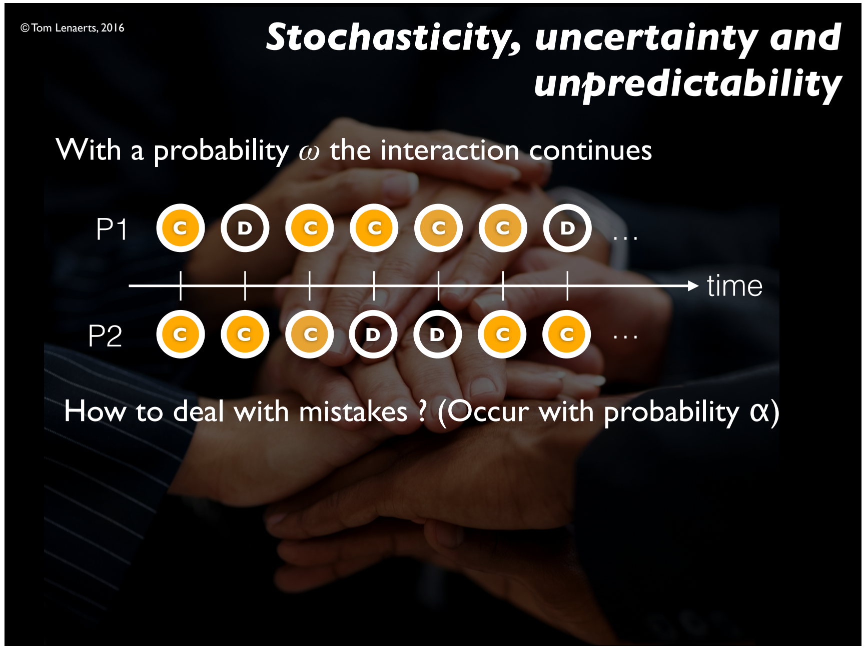

58-1

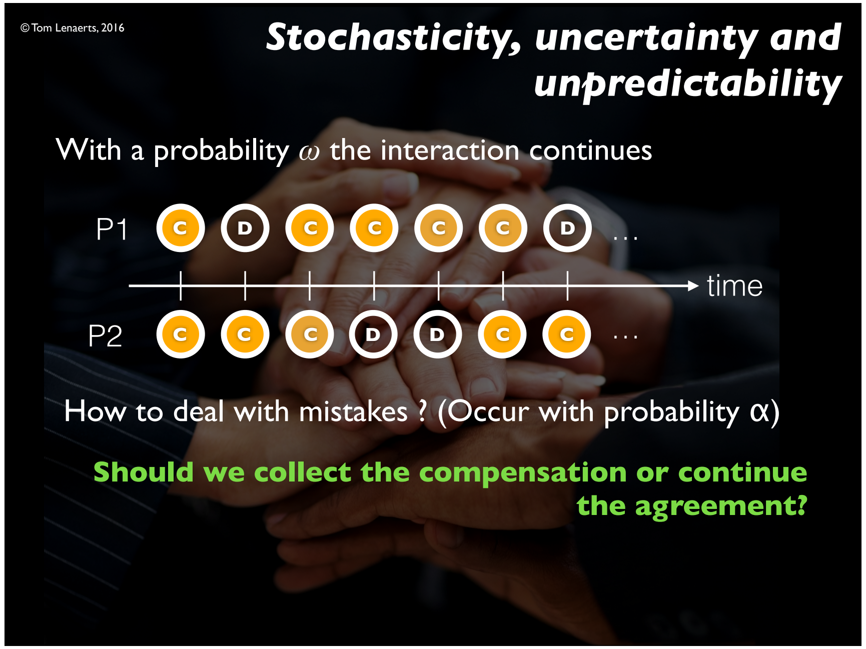

58-2

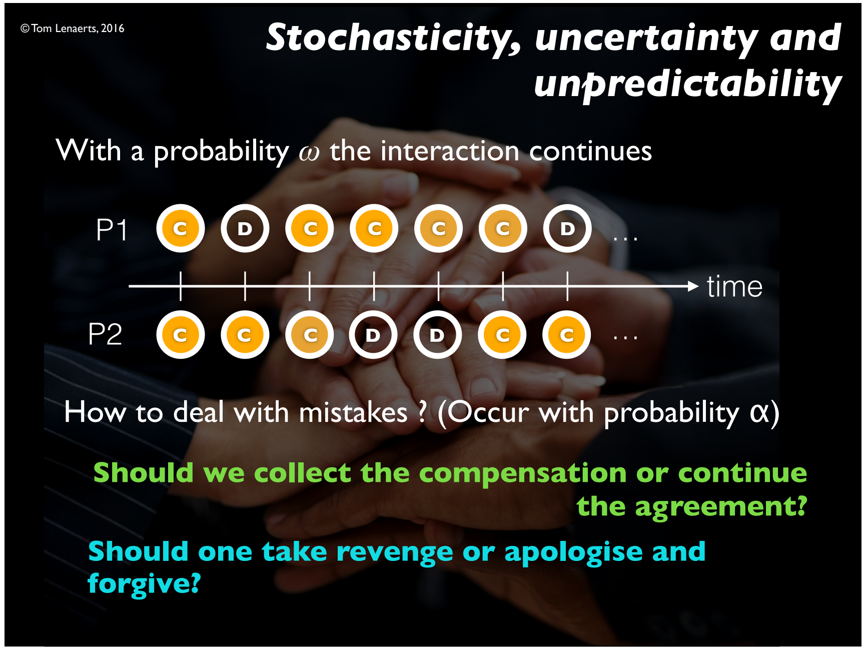

58-3

59

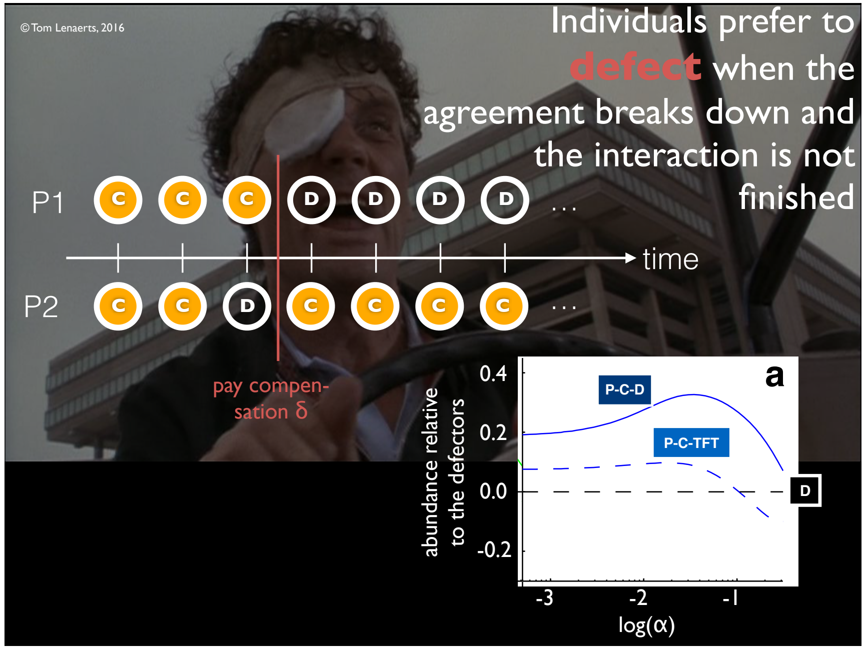
©Tom Lenaerts, 2016
Individuals prefer to defect when the agreement breaks down and the interaction is not finished
P1
P2
time
pay compen-sation δ
abundance relative to the defectors
P-C-D
P-C-TFT
a
D
0.4
0.2
0.0
-0.2
-3
-2
-1
log(α)

60

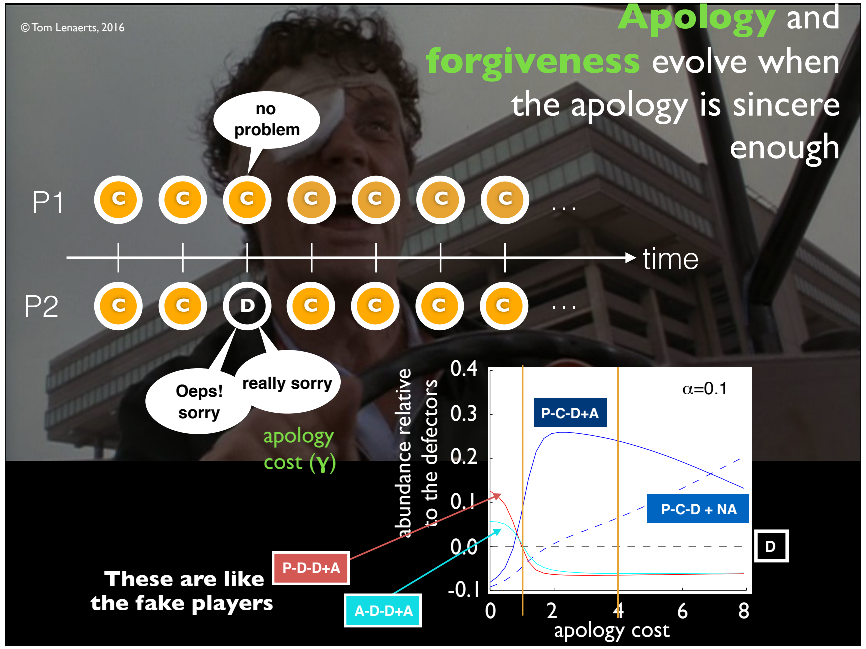
©Tom Lenaerts, 2016
Apology and forgiveness evolve when the apology is sincere enough
no problem
P1
P2
time
Oeps! sorry
really sorry
apology cost (γ)
These are like the fake players
P-D-D+A
A-D-D+A
abundance relative to the defectors
α=0.1
P-C-D+A
P-C-D + NA
D
0.4
0.3
0.2
0.1
0.0
-0.1
0
2
4
6
8
apology cost

61


©Tom Lenaerts, 2016
How do people really behave in games?

62

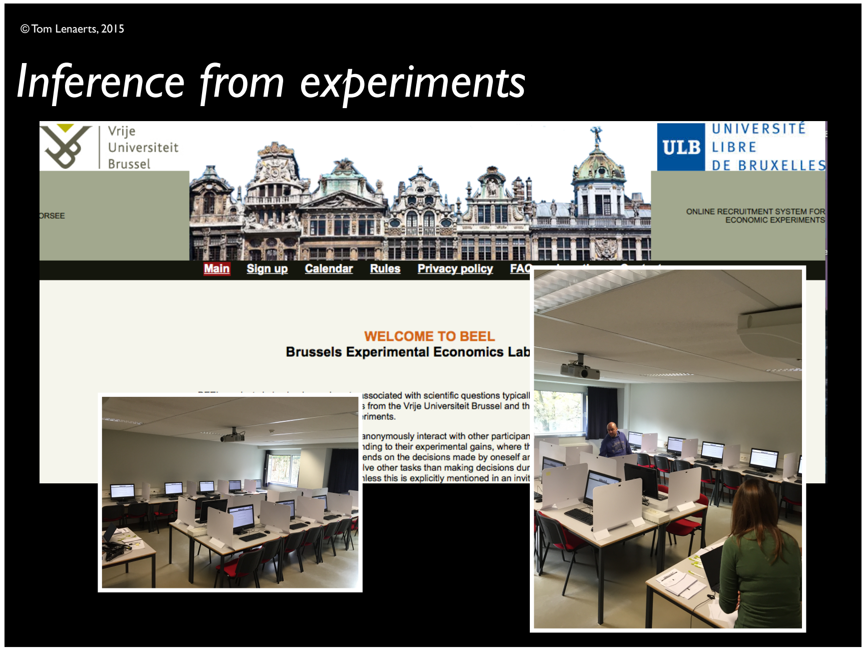
©Tom Lenaerts, 2015
Inference from experiments
Vrije Universiteit Brussel
ULB UNIVERSITÉ LIBRE DE BRUXELLES
ONLINE RECRUITMENT SYSTEM FOR ECONOMIC EXPERIMENTS
Main
Sign up
Calendar
Rules
Privacy policy
WELCOME TO BEEL
Brussels Experimental Economics Lab

63

Tom Lenaerts - Omina 2019

How to arrive at a legally binding universal agreement to avoid a global 2° temperature increase?

64

Tom Lenaerts - Omina 2019

Voluntary pledges require **trust** and **cooperative** participants

Current Climate Pledges Aren't Enough to Stop Severe Warming

Existing pledges under the Paris Agreement won't curb warming to two degrees Celsius above preindustrial levels, a new UN report warns.

3 MINUTE READ

BY MICHAEL GRESHKO

PUBLISHED OCTOBER 31, 2017

The world is not doing enough to curb its collective carbon emissions, a new UN report warns.

In an audit of the global Paris Agreement released Tuesday, the UN Environment Programme finds that if action to combat climate change is limited to just current pledges, the Earth will get at least three degrees Celsius (5.4 degrees Fahrenheit) warmer by 2100 relative to preindustrial levels.

This amount of warming would vastly exceed the Paris Agreement's goal, which is to limit global warming by the end of the century to two degrees

SCIENTIFIC REPORTS

OPEN Good Agreements Make Good Friends

The Anh Han[1,2], Luís Moniz Pereira[3], Francisco C. Santos[4,5] & Tom Lenaerts[1,2]

Non-costly, non-enforced agreements lead to defection

65

Tom Lenaerts - Omina 2019

How to study decision-making in the climate change issue.

66-1

Tom Lenaerts - Omina 2019

How to study decision-making in the climate change issue.

How to model this problem?

66-2

Tom Lenaerts - Omina 2019

How to study decision-making in the climate change issue.

How to model this problem?

How do humans make decisions?

66-3

Tom Lenaerts - Omina 2019

How to study decision-making in the climate change issue.

How to model this problem?

How do humans make decisions?

What features influence this decision?

66-4

Tom Lenaerts - Omina 2019

Uncertainty encourages group reciprocation and polarization in a high-risk climate change problem

Elias Fernández Domingos[1,2,6,*], Jelena Grujić[1,2,*], Juan C. Burguillo[6], Georg Kirchsteiger[3], Francisco C. Santos[4,5], and Tom Lenaerts[1,2,‡]

[1]AI lab, Computer Science Department, Vrije Universiteit Brussel, Pleinlaa 9, 3rd floor, 1050 Brussels, Belgium

[2]MLG, Département d'Informatique, Université Libre de Bruxelles, Boulevard du Triomphe, CP 212, 1050 Brussels, Belgium

[3]ECARES, Université Libre de Bruxelles, Av. Roosevelt 42, CP 114, 1050 Brussels, Belgium

[4]INESC-ID and Instituto Superior Técnico, Universidade de Lisboa, IST-Taguspark, 2744-016 Porto Salvo, Portugal

[5]ATP-group, 2744-016 Porto Salvo, Portugal.

[6]Department of Telematic Engineering, University of Vigo, 36310 Vigo, Spain

PRELIMINARY

67

Tom Lenaerts - Omina 2019

The game : Collective risk dilemma

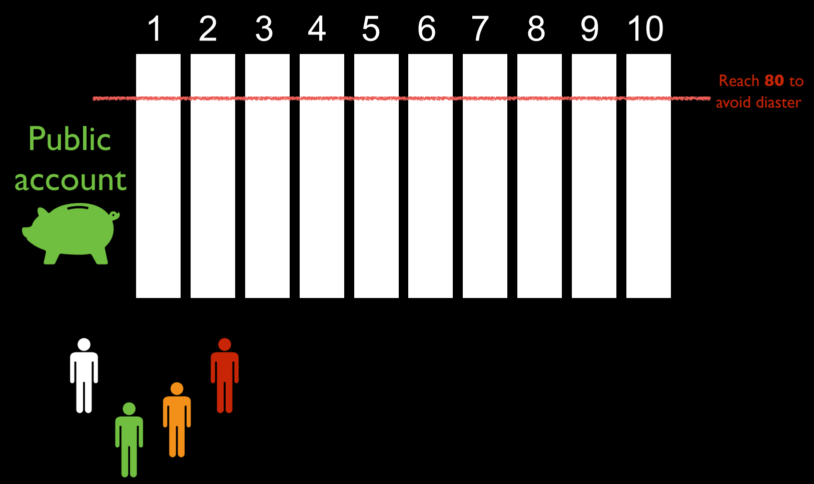

68-1

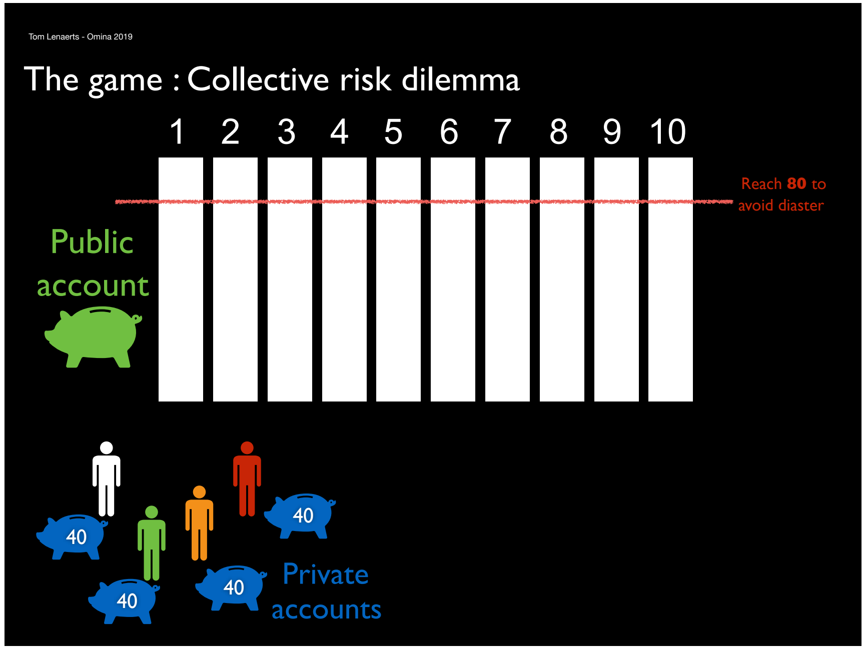

68-2

68-3

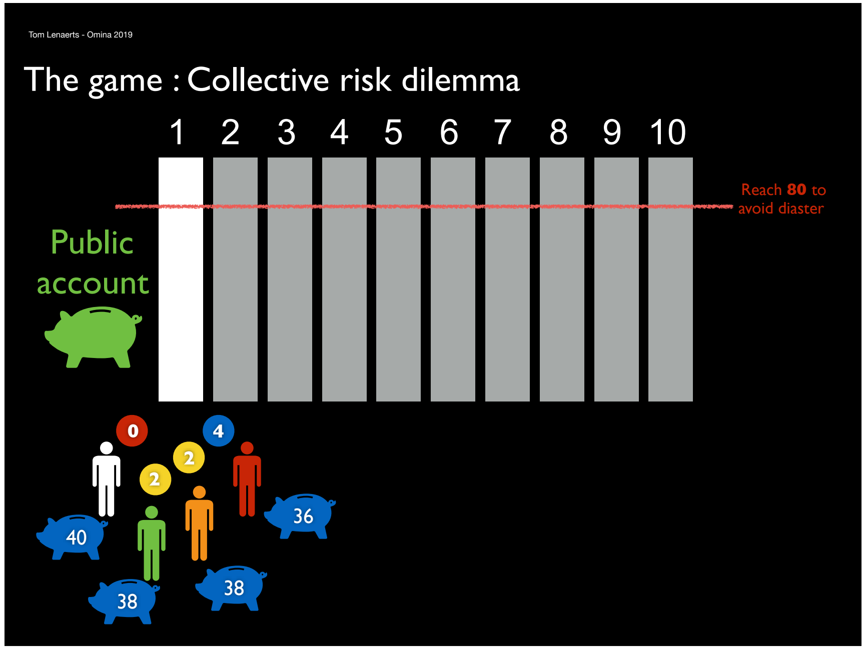

69-1

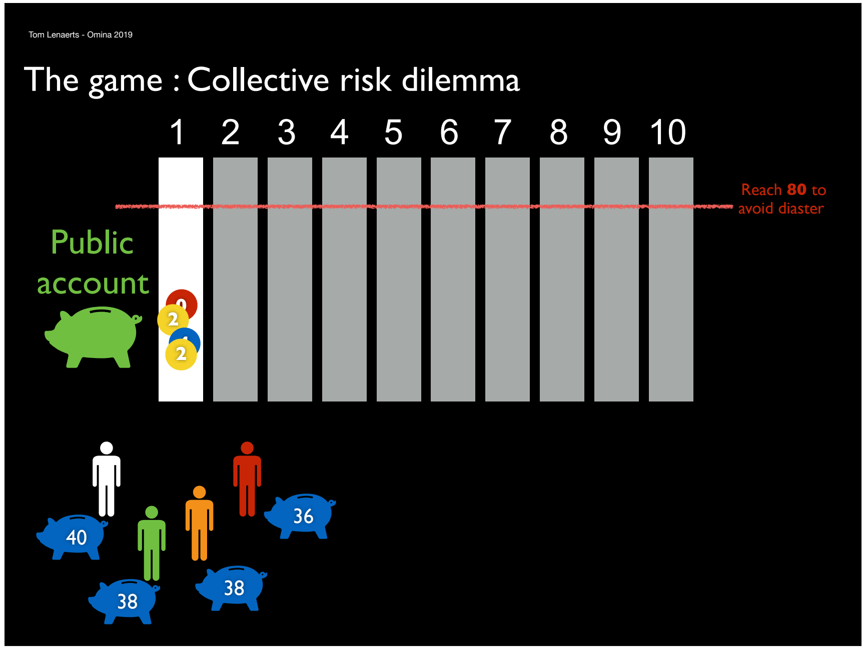

69-2

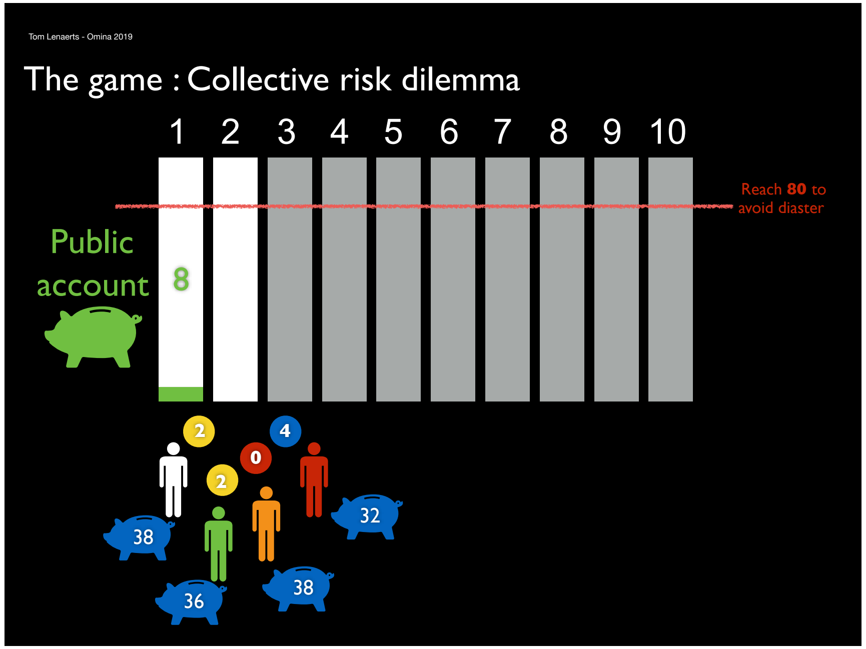

70-1

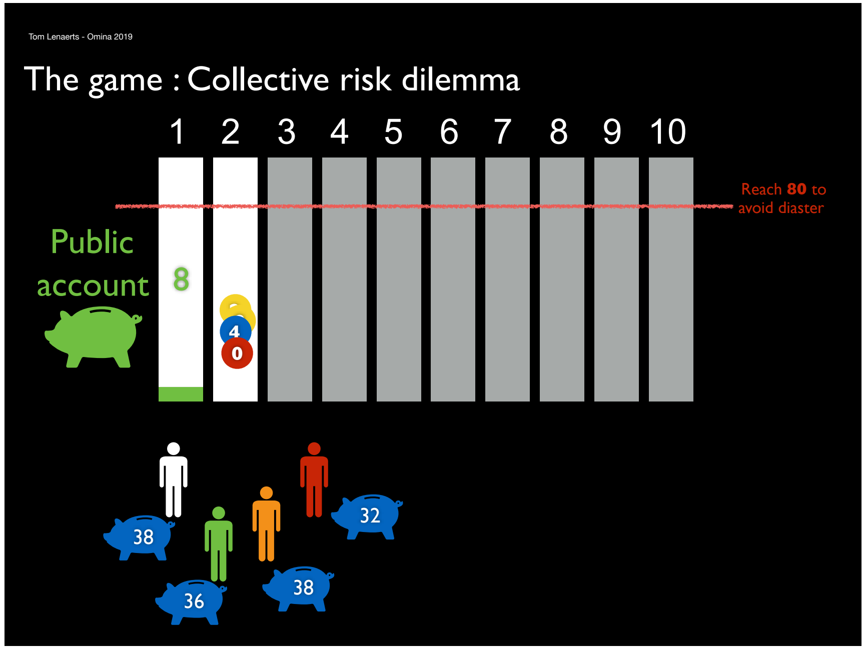

70-2

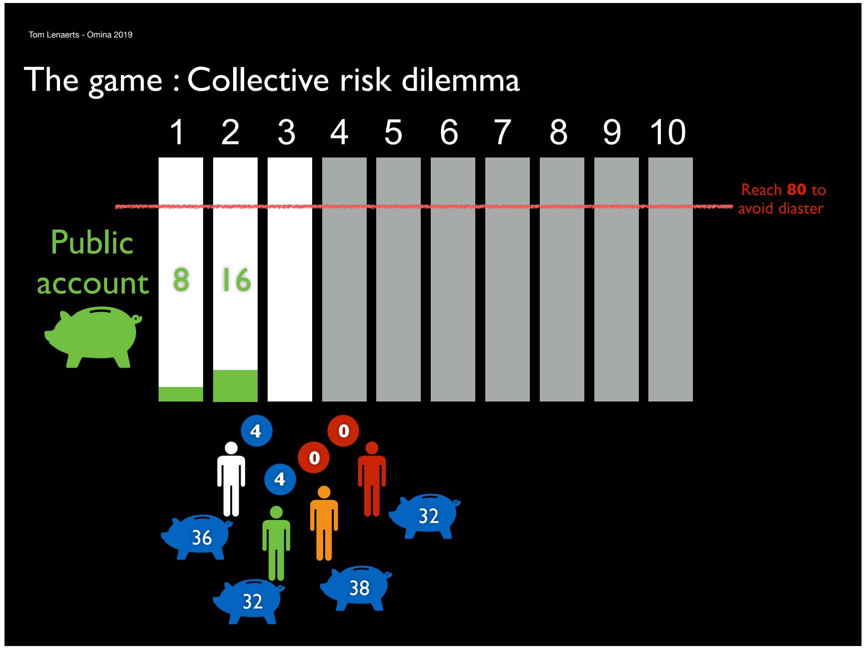

71-1

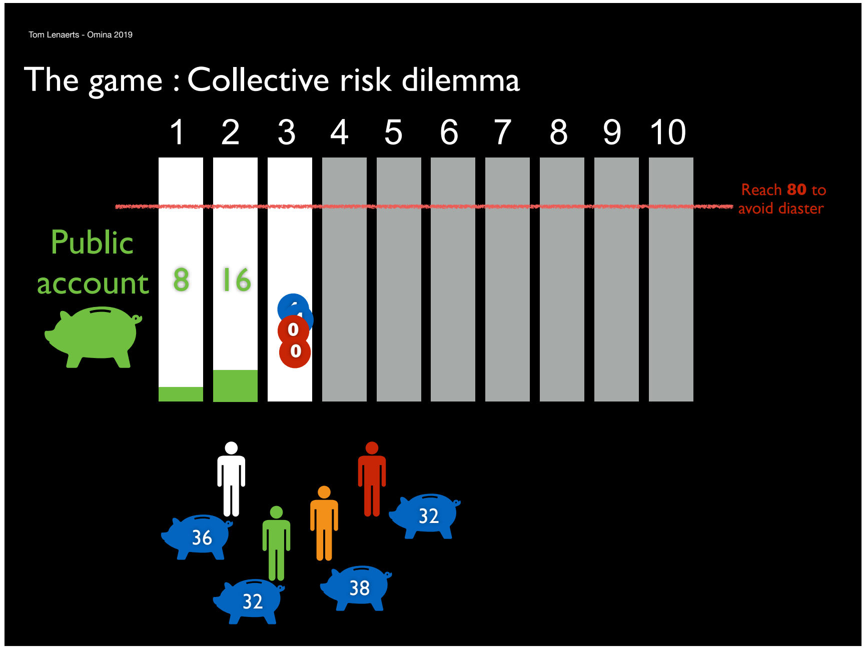

71-2

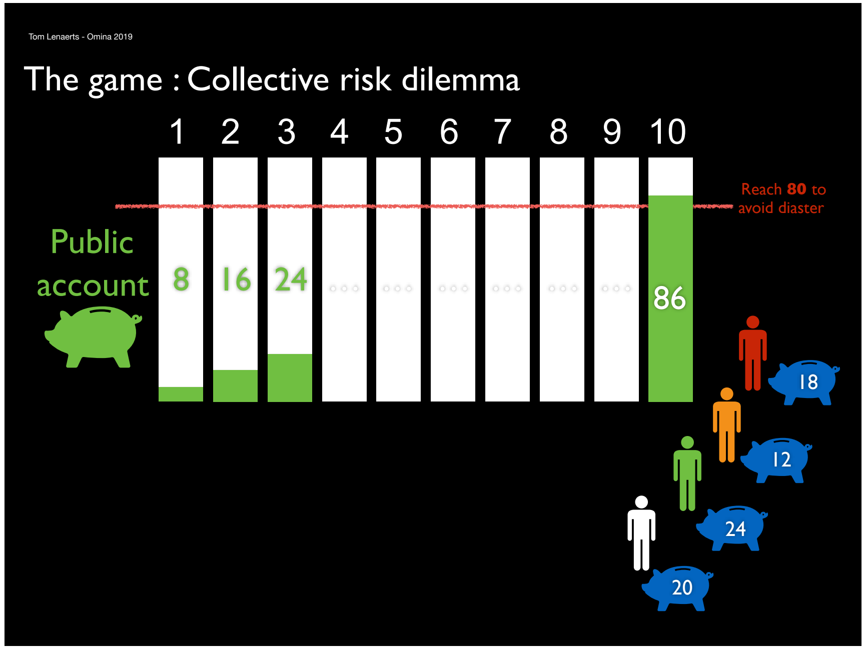

72

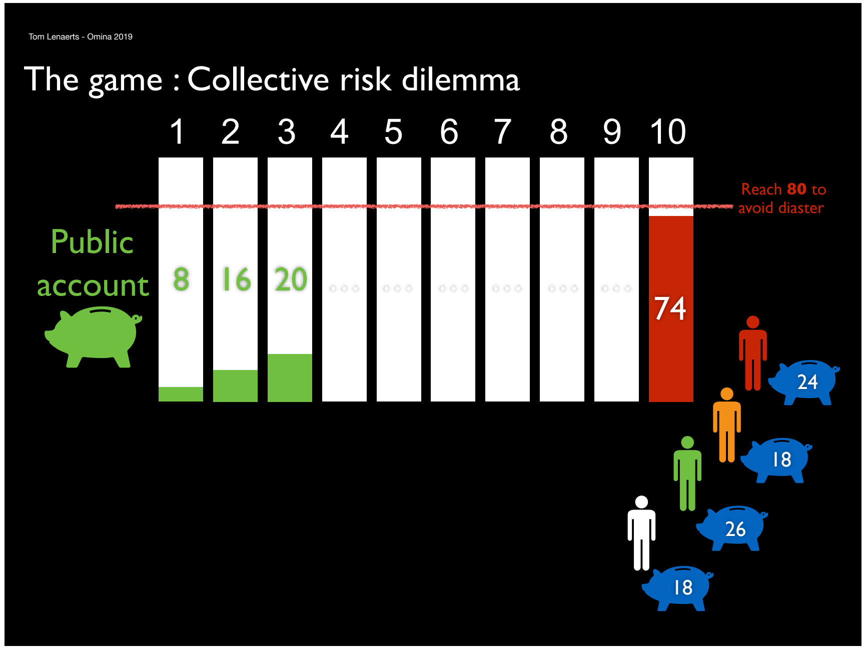

73-1

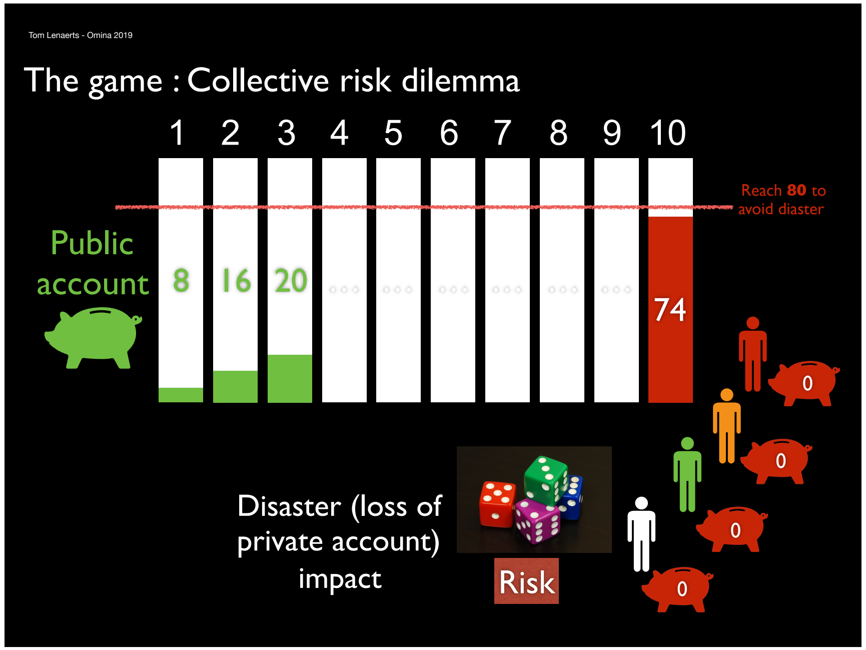

73-2

74-1

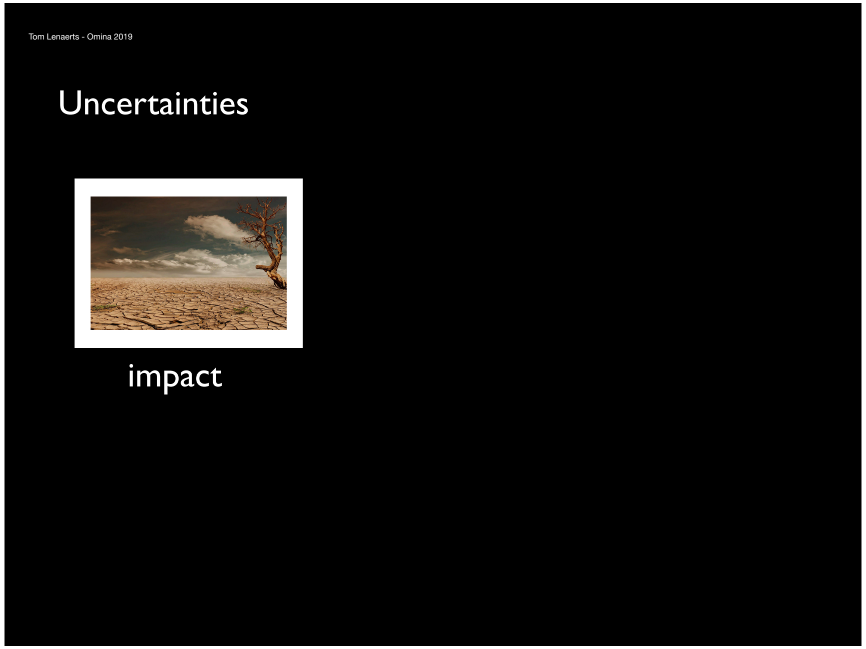

74-2

74-3

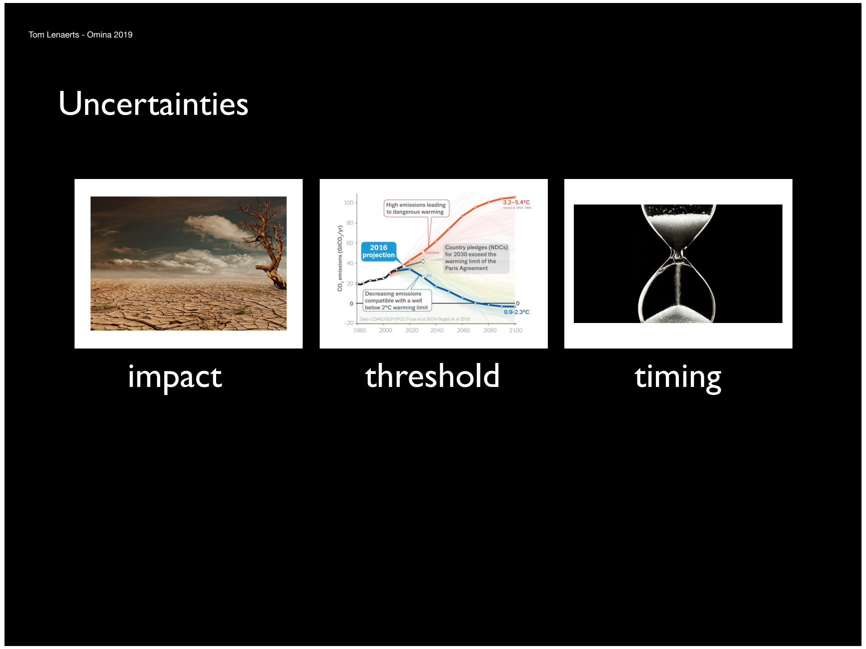

74-4

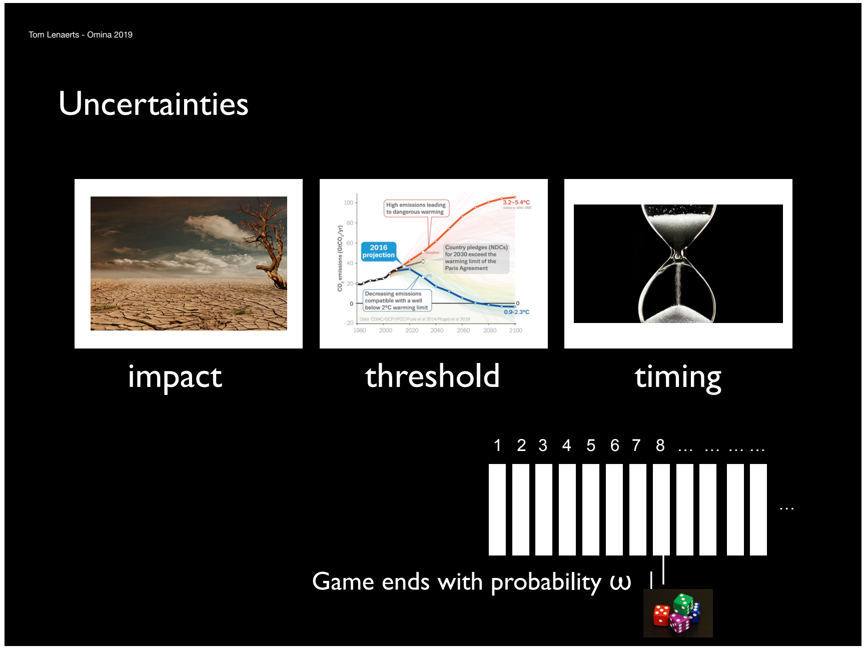

74-5

Tom Lenaerts - Omina 2019

No uncertainty

6 participants

Risk 90%

Threshold 120

Endowment 40

Rounds 10

75-1

Tom Lenaerts - Omina 2019

No uncertainty

6 participants

Risk 90%

Threshold 120

Endowment 40

Rounds 10

Low uncertainty

Same **BUT**

ω =1/3
min. 8 rounds

Avg. rounds 10

75-2

Tom Lenaerts - Omina 2019

No uncertainty

6 participants

Risk 90%

Threshold 120

Endowment 40

Rounds 10

Low uncertainty

Same **BUT**

ω =1/3
min. 8 rounds

Avg. rounds 10

High uncertainty

Same **BUT**

ω =1/5
min. 6 rounds

Avg. rounds 10

75-3

Tom Lenaerts - Omina 2019

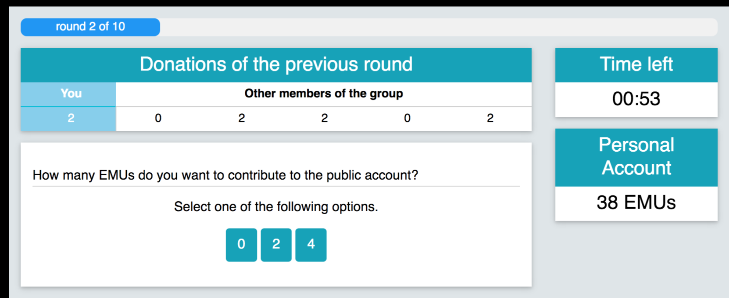

76

Tom Lenaerts - Omina 2019

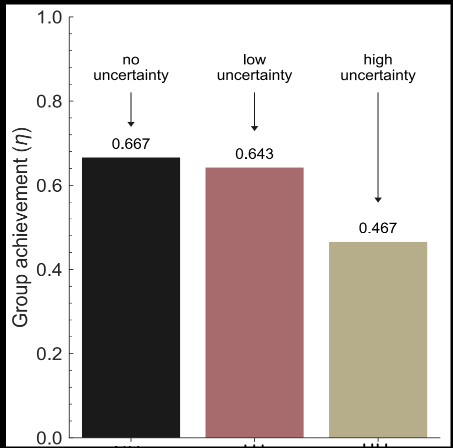

No significant differences in group success

77

Tom Lenaerts - Omina 2019

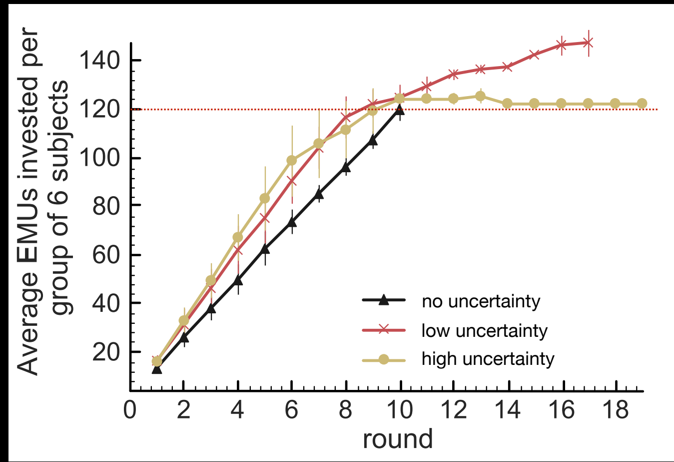

Players invest early in when there is timing uncertainty

78-1

Tom Lenaerts - Omina 2019

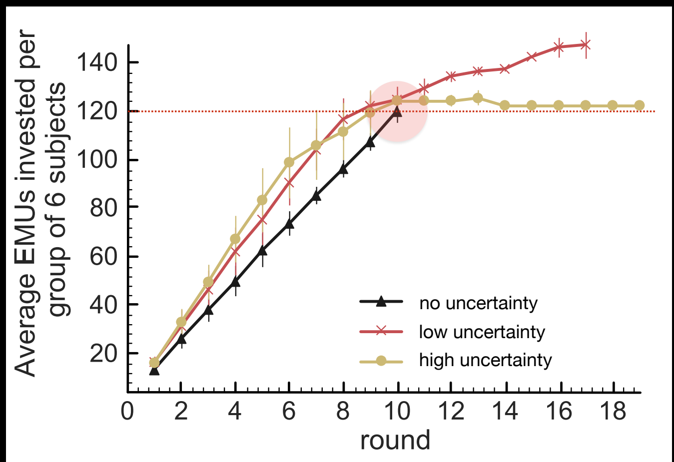

Players invest early in when there is timing uncertainty

78-2

Tom Lenaerts - Omina 2019

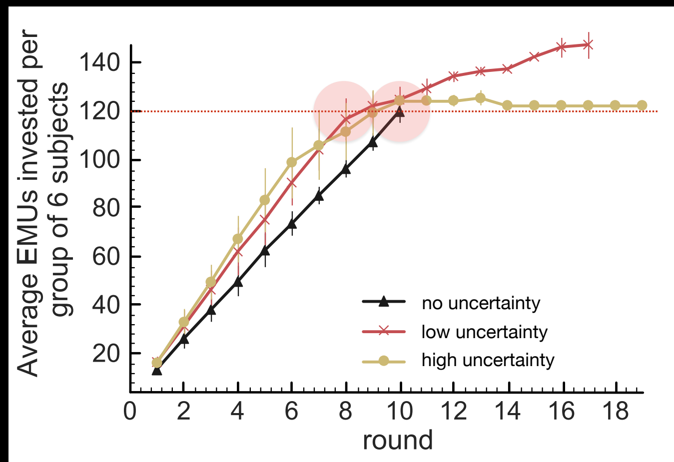

Players invest early in when there is timing uncertainty

78-3

Tom Lenaerts - Omina 2019

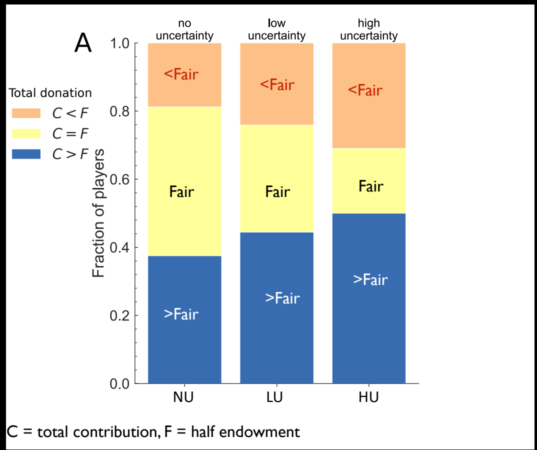

Participants become more polarised

79

Tom Lenaerts - Omina 2019

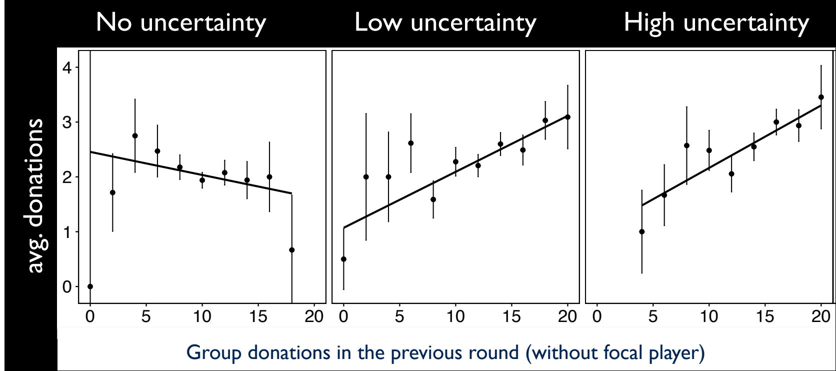

Group reciprocal strategies are needed to avoid disaster

80

Tom Lenaerts - Omina 2019

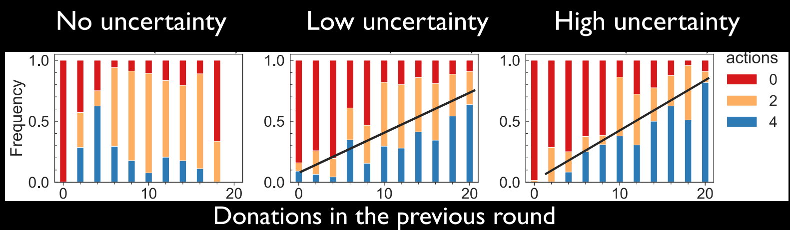

What other behaviours ?

81-1

Tom Lenaerts - Omina 2019

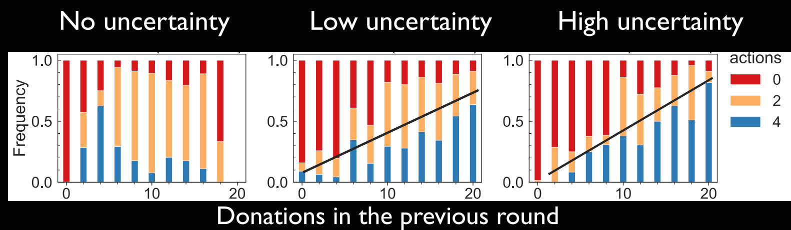

Uncertainty induces reciprocal behaviour

What other behaviours ?

81-2

Tom Lenaerts - Omina 2019

F **Fair** players, give 2

D Defecting players (**free riders**), give 0

A **Altruistic** players, give 4

C **Compensators** players, give 0 when prior investment was >=10, otherwise 4

R **Reciprocators** players, give 4 when prior investment was >=10, otherwise 0

*All end when threshold is reached

82

Tom Lenaerts - Omina 2019

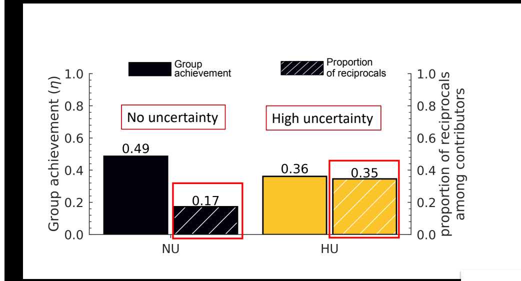

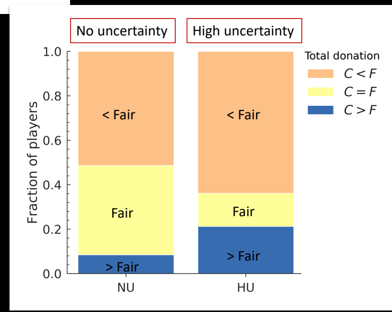

83

Tom Lenaerts - Omina 2019

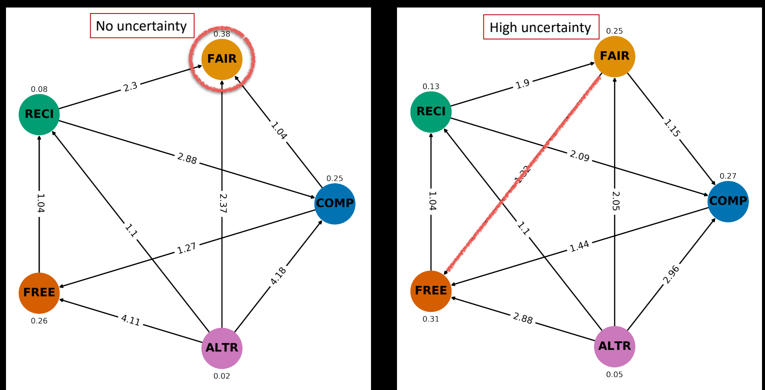

84

Uncertainty influences behaviour

Timing uncertainty

Leads to **polarisation**

Success requires **reciprocal** mechanisms

Investing early promotes coordination among peers

EGT models explain experiments

85

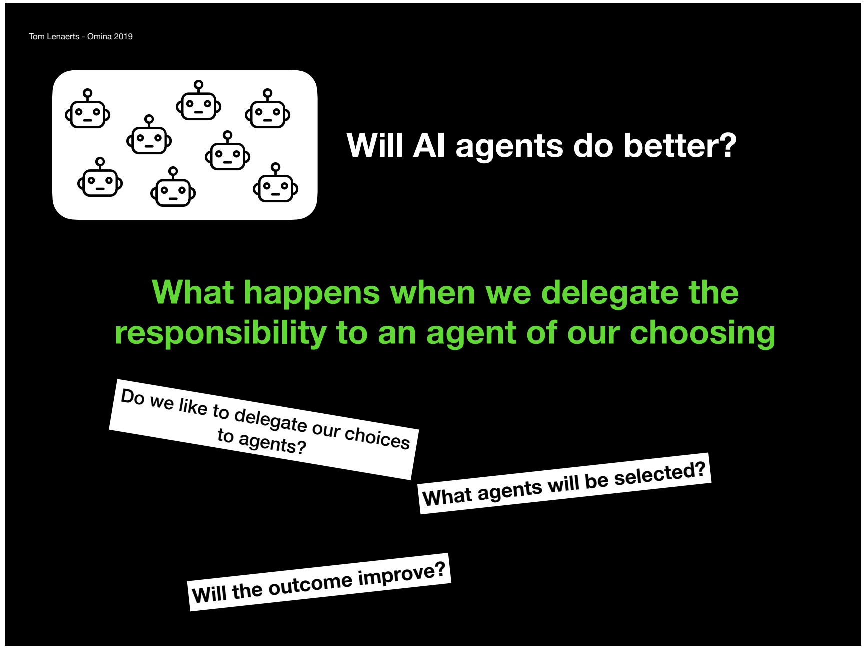
Tom Lenaerts - Omina 2019
Will AI agents do better?
What happens when we delegate the responsibility to an agent of our choosing
Do we like to delegate our choices to agents?
What agents will be selected?
Will the outcome improve?

86

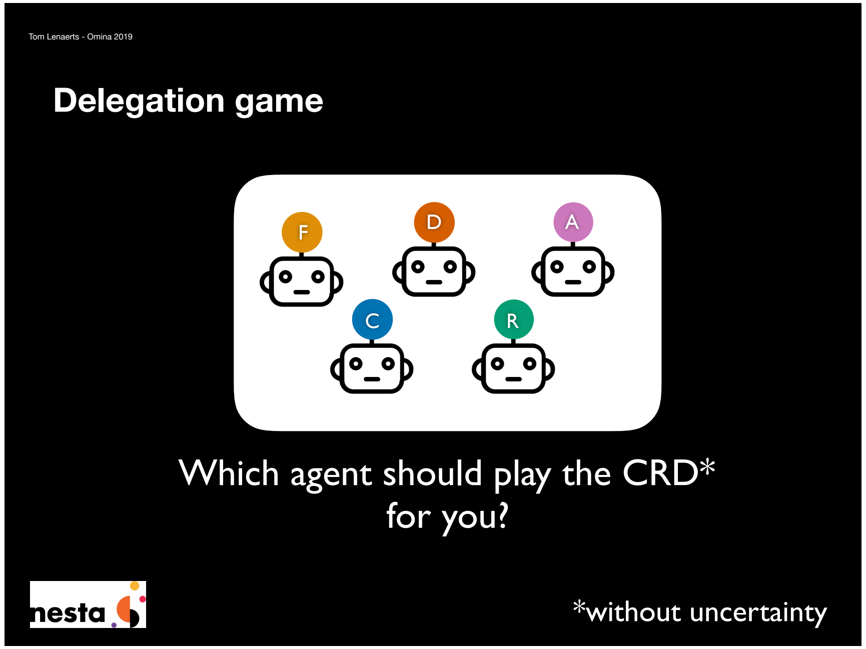
Tom Lenaerts - Omina 2019
Delegation game
F
D
A
C
R
Which agent should play the CRD* for you?
nesta
*without uncertainty

87

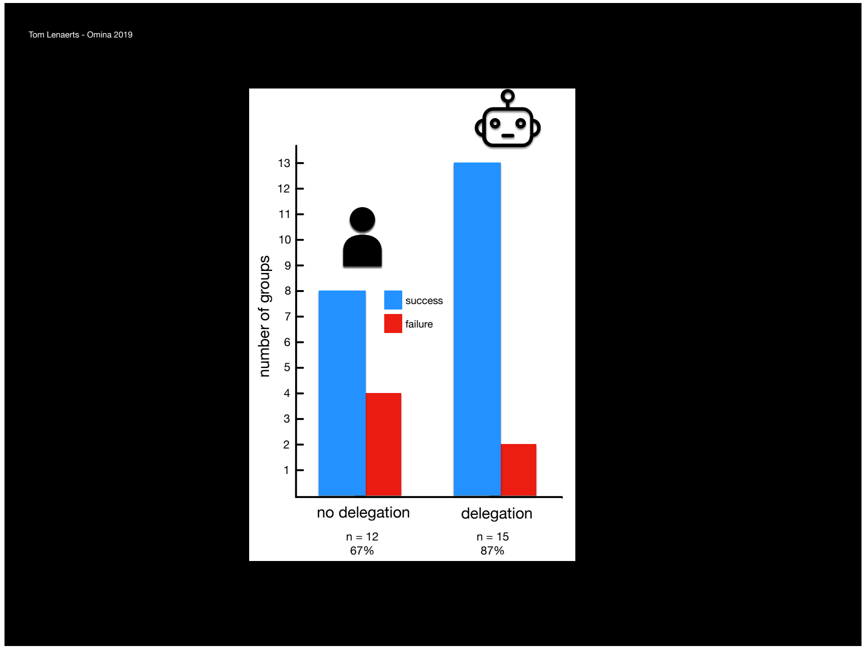
Tom Lenaerts - Omina 2019
number of groups
13
12
11
10
9
8
7
6
5
4
3
2
1
success
failure
no delegation
n = 12
67%
delegation
n = 15
87%

88

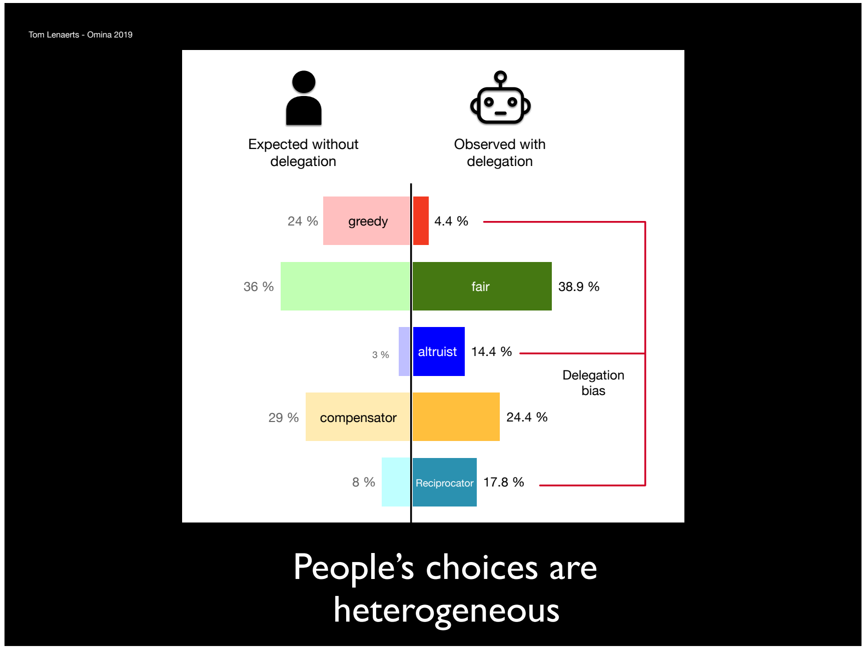
Tom Lenaerts - Omina 2019
Expected without delegation
Observed with delegation
24 %
greedy
4.4 %
36 %
fair
38.9 %
3 %
altruist
14.4 %
Delegation bias
29 %
compensator
24.4 %
8 %
Reciprocator
17.8 %
People's choices are heterogeneous

89

Tom Lenaerts - Omina 2019

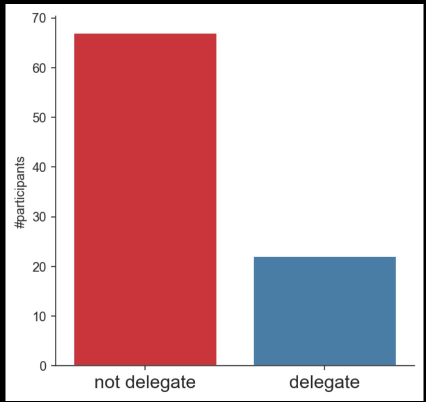

They would not delegate again

90

Tom Lenaerts - Omina 2019

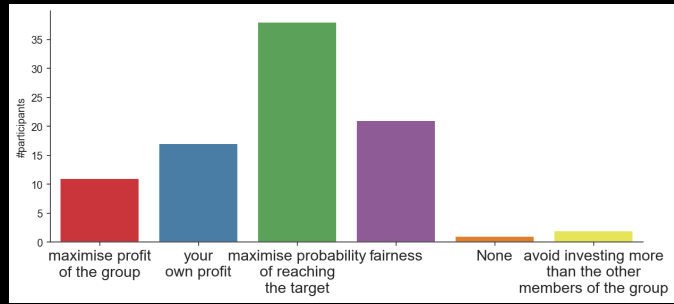

Most wanted to reach the target

91

Tom Lenaerts - Omina 2019

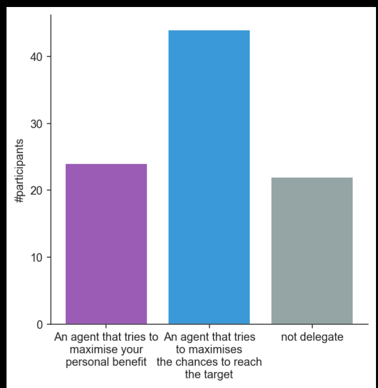

Would select and agent that maximises the chance to reach the target

92

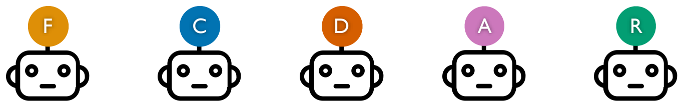

Delegation to agents increases success, but **humans prefer to act for themselves**

All understood that the goal is to reach the target to guarantee a benefit

Tom Lenaerts - Omina 2019

93

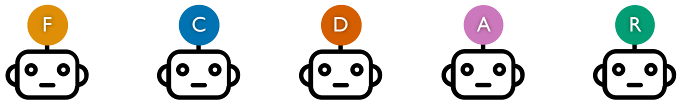

Delegation forces to **think about future rewards**

Fear of betrayal in the game disappears as behaviour is fixed

Social norms are no longer used when delegating to an agent

Aimone, J. A., & Houser, D. (2012). What you don't know won't hurt you: a laboratory analysis of betrayal aversion. *Experimental Economics*, *15*(4), 571-588.
Ensthaler, L., Huck, S., & Leutgeb, J. (2019). Games played through agents in the laboratory—a test of Prat & Rustichini's model. *Games and Economic Behavior*.
Prat, A., & Rustichini, A. (2003). Games played through agents. *Econometrica*, *71*(4), 989-1026.

Questions