While you are waiting...

- **socrative.com,** room number **SIMLANG2016**

Simulating Language
Lecture 8: Learning bias considered

Simon Kirby

simon@ling.ed.ac.uk

Rule: [1, -1, -1, 0]

Passes acquisition test? **Yes**

Maintenance: **Yes**

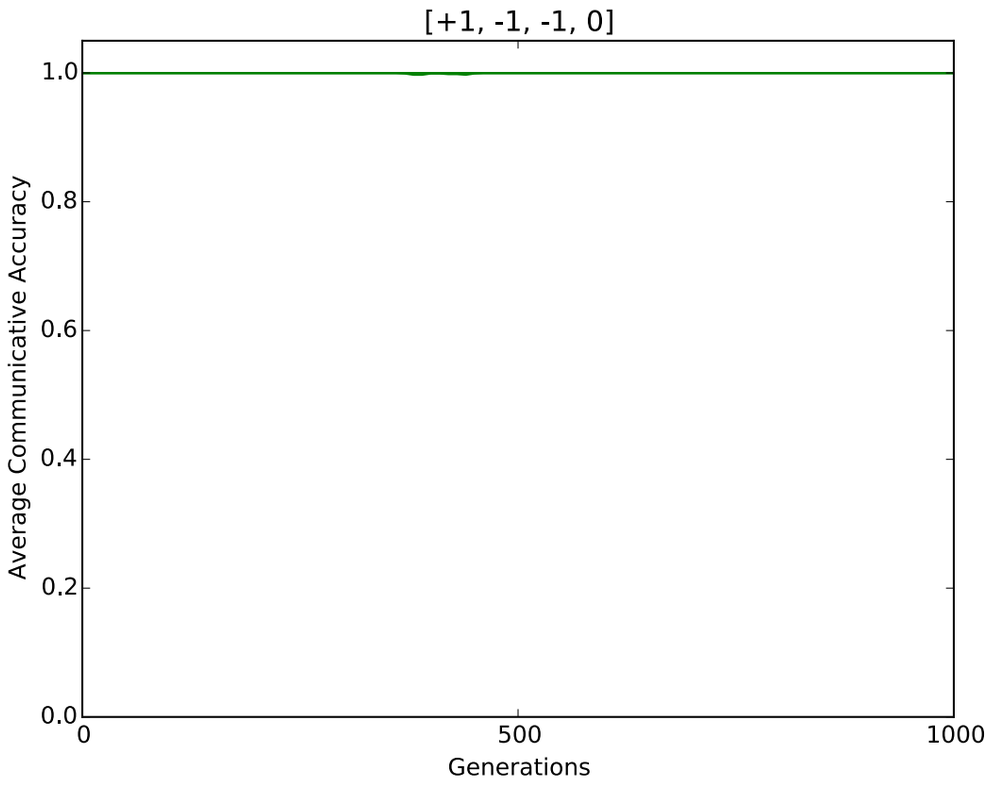

Construction: **Yes**

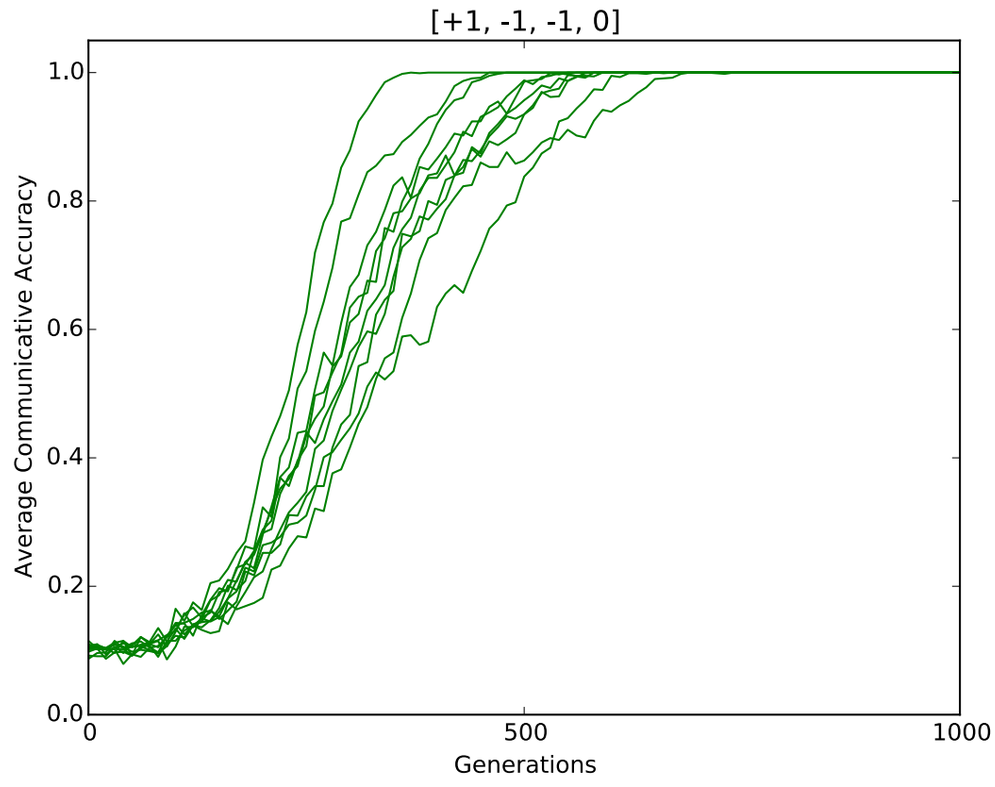

Rule: [1, 0, 0, 0]

Passes acquisition test? **Yes**

Maintenance: **Yes**

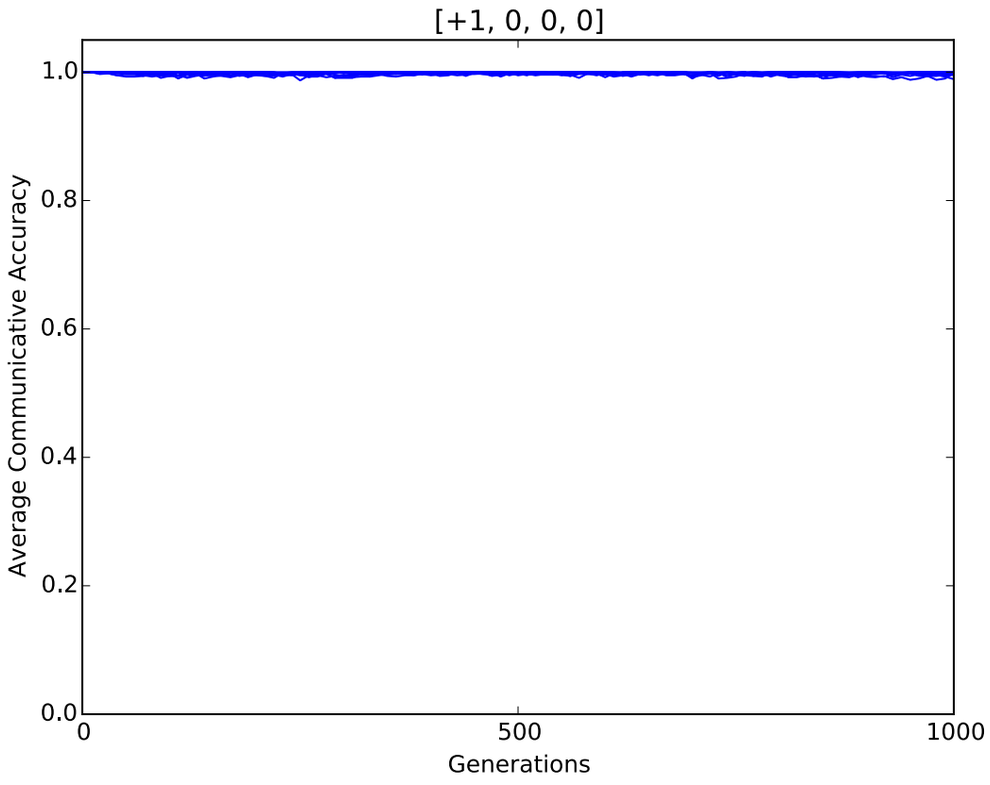

Construction: **No**

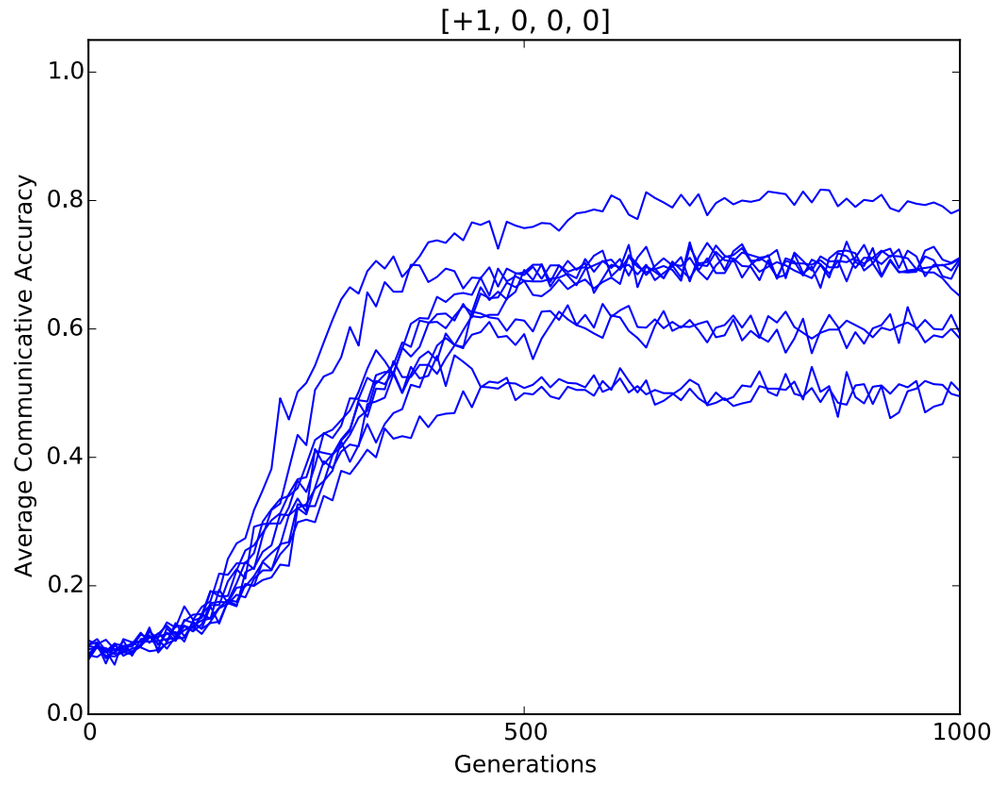

Rule: [1, -1, 1, 0]

Passes acquisition test? **Yes**

Maintenance: **No**

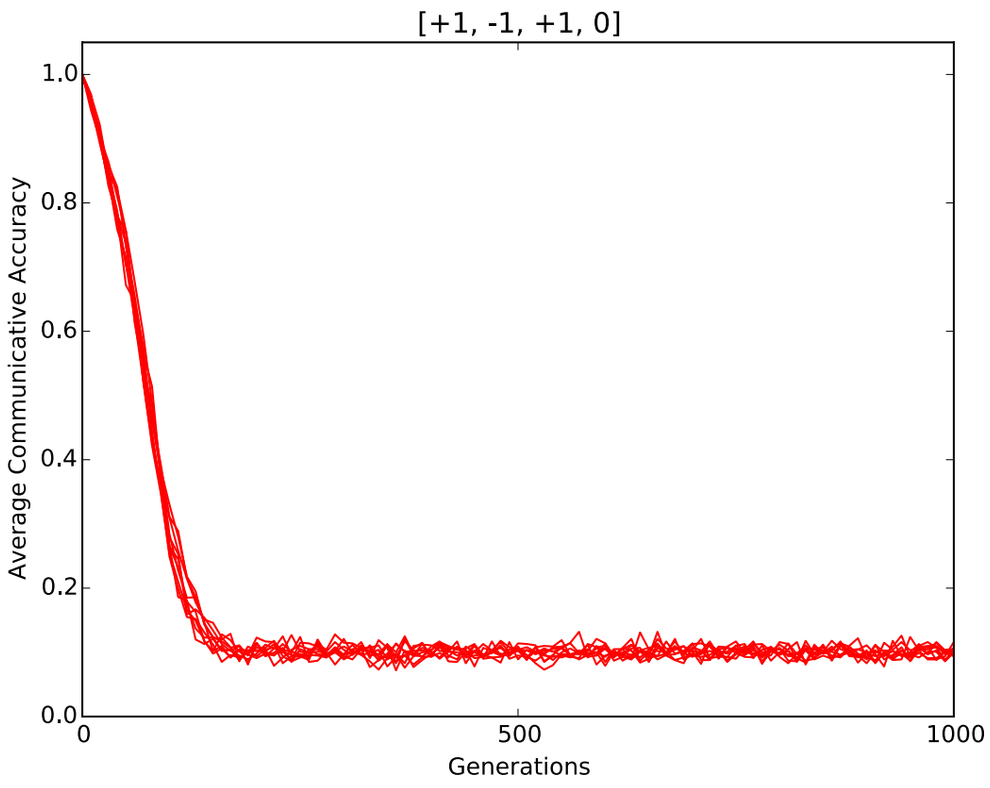

Construction: **No**

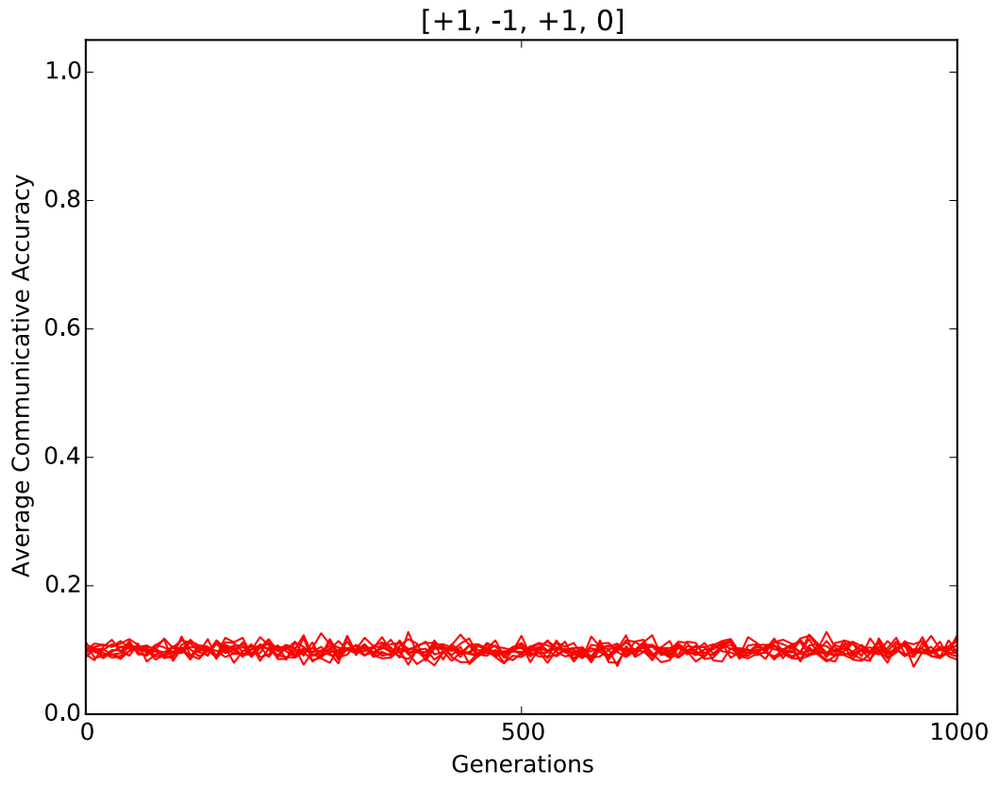

Bias

- Different weight update rules correspond to different ways of learning
- They come with different *biases*
    - Although that’s not immediately obvious just from looking at acquisition
- Population’s language (in this case, just a vocabulary really) evolves to fit these biases
- Biases are a consequence of α, β, γ and δ
- But what exactly are these different biases?

Working out bias

- A constructor rule: [+1, -1, -1, +1]

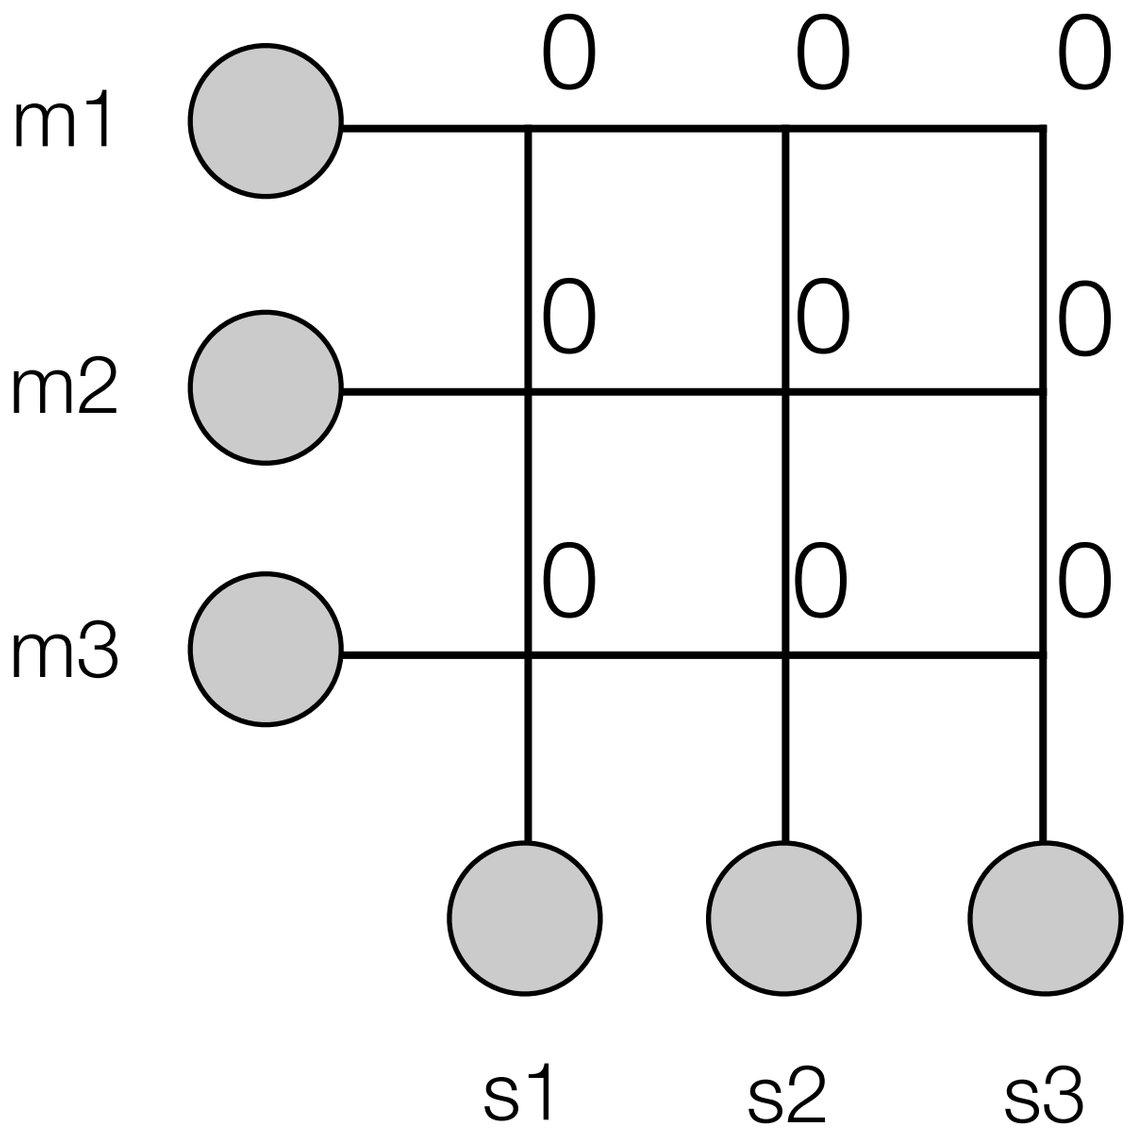

Working out bias

- A constructor rule: [+1, -1, -1, +1]

Observation:
m1→s1

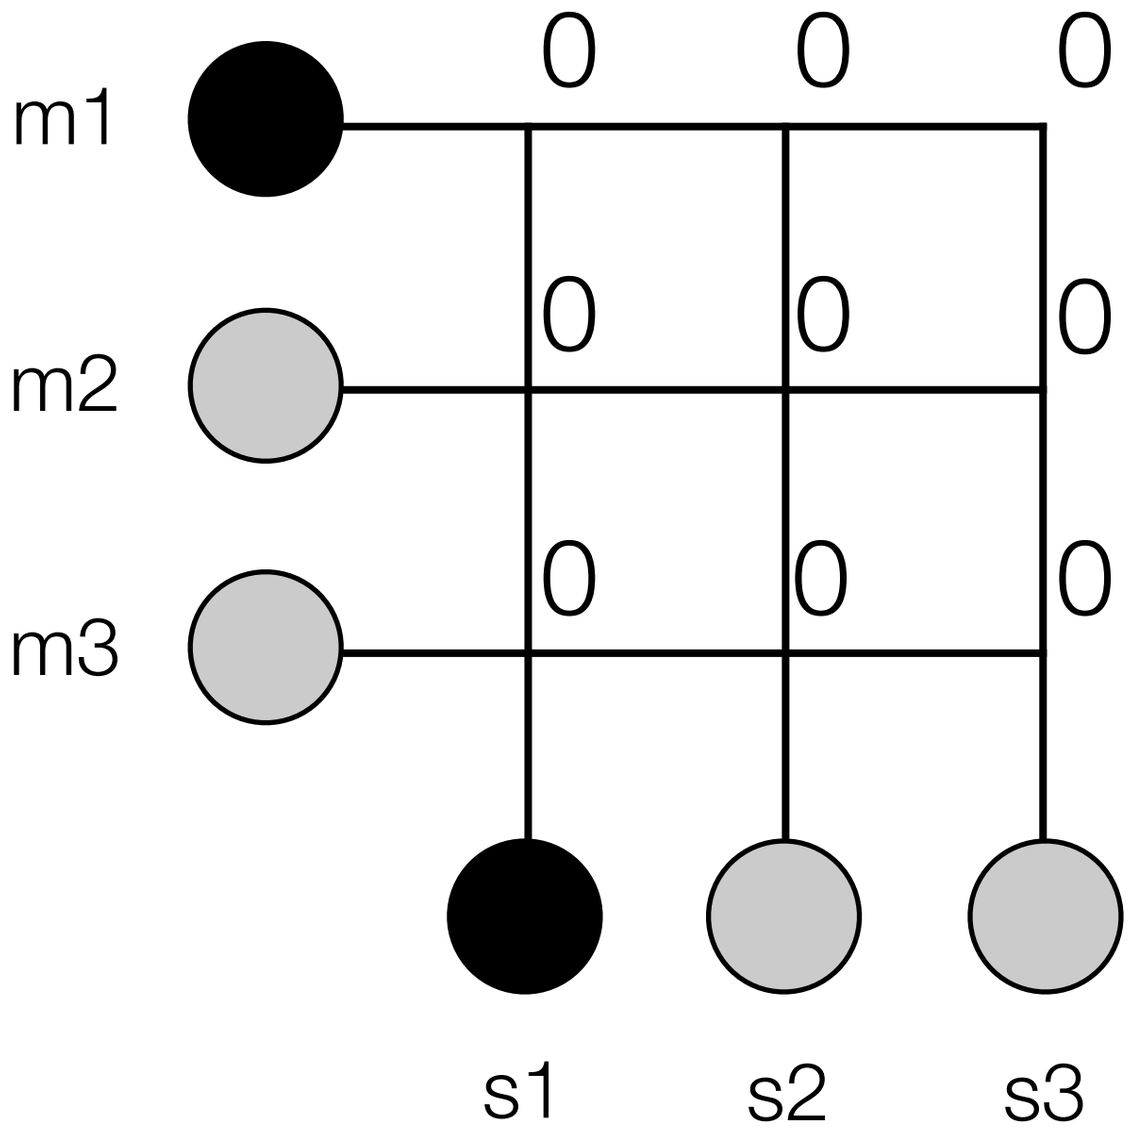

Working out bias

- A constructor rule: [+1, -1, -1, +1]

Observation:
m1→s1

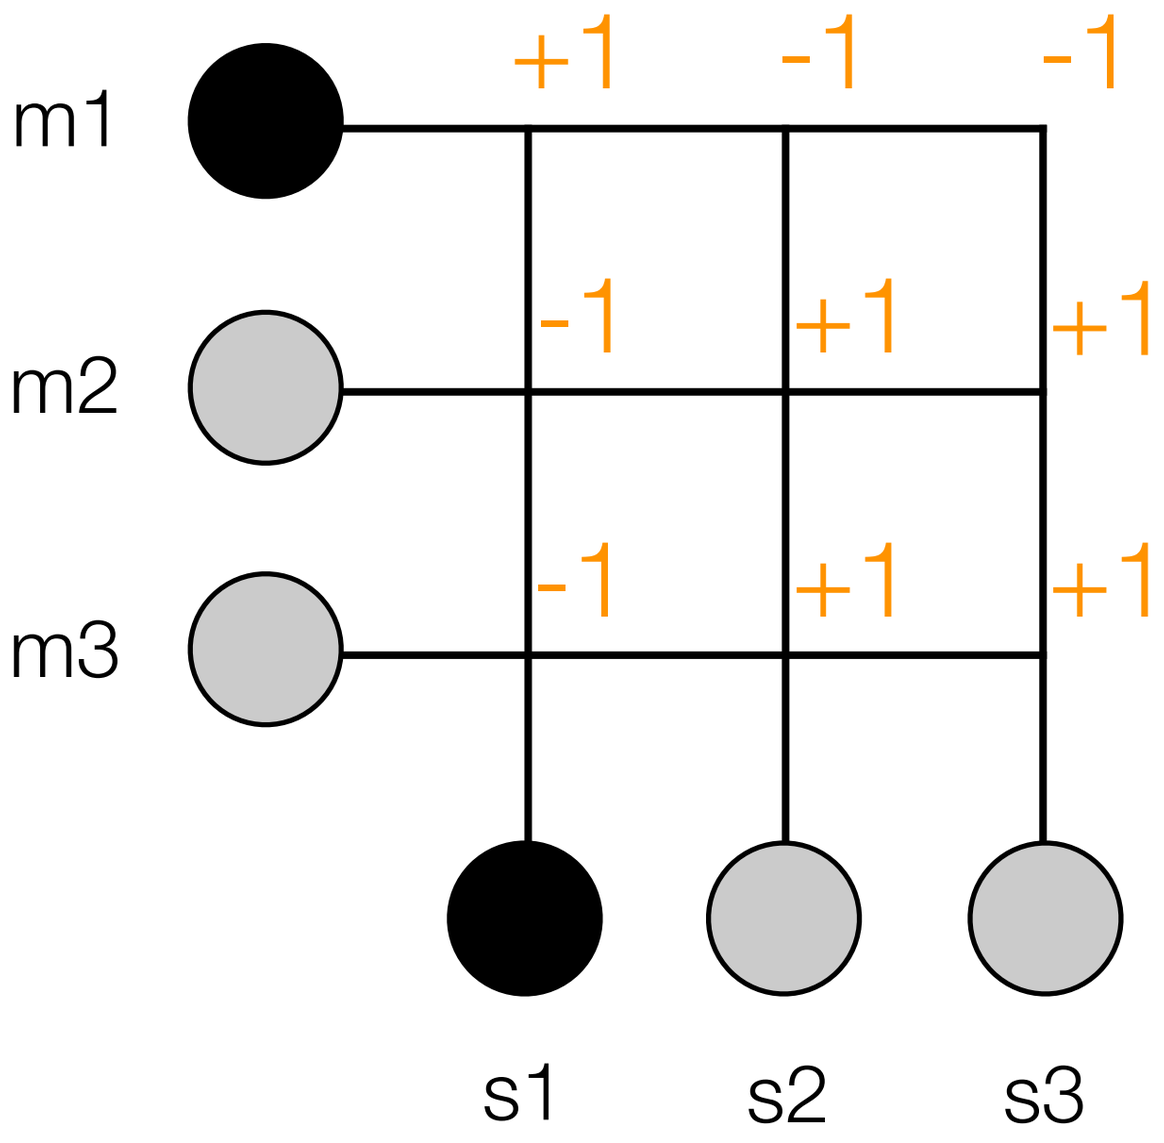

Working out bias

- A constructor rule: [+1, -1, -1, +1]

Observation:
m1→s1

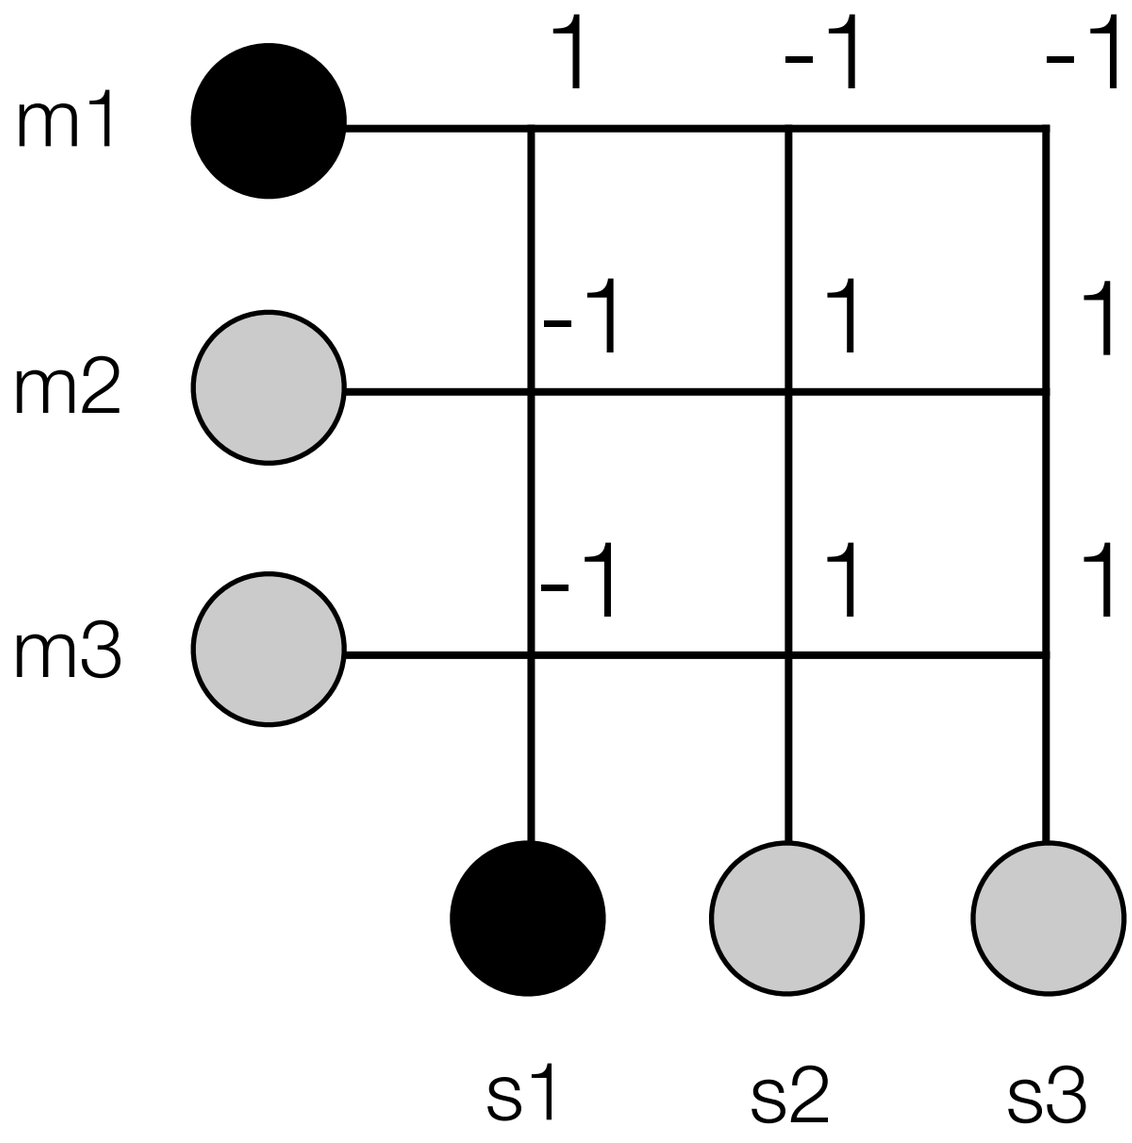

Working out bias

- A constructor rule: [+1, -1, -1, +1]

Production:

m1 → ?

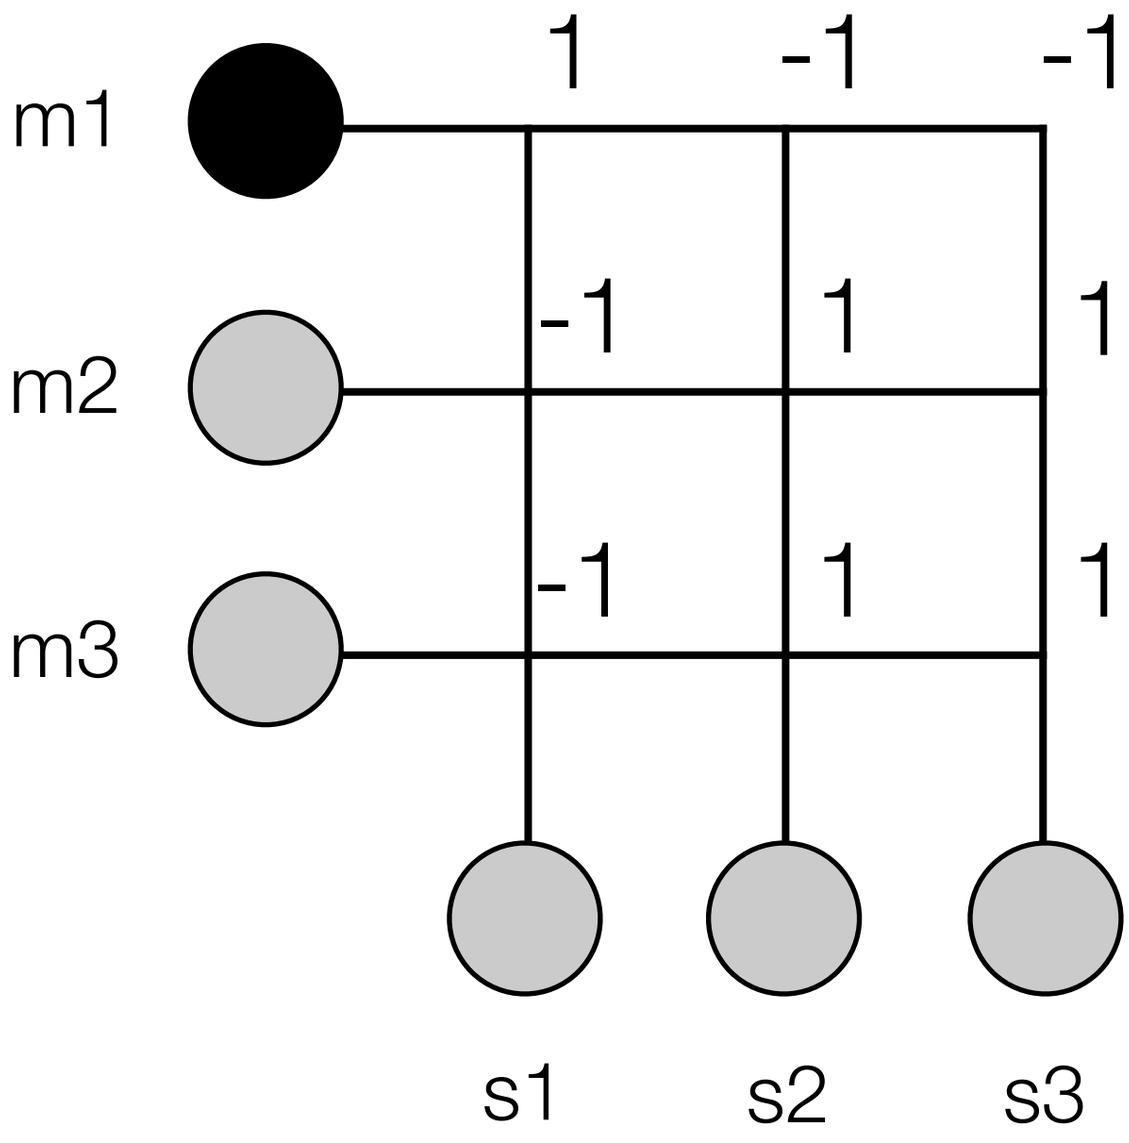

Working out bias

- A constructor rule: [+1, -1, -1, +1]

Production:

m1→s1

(*not* s2 or s3)

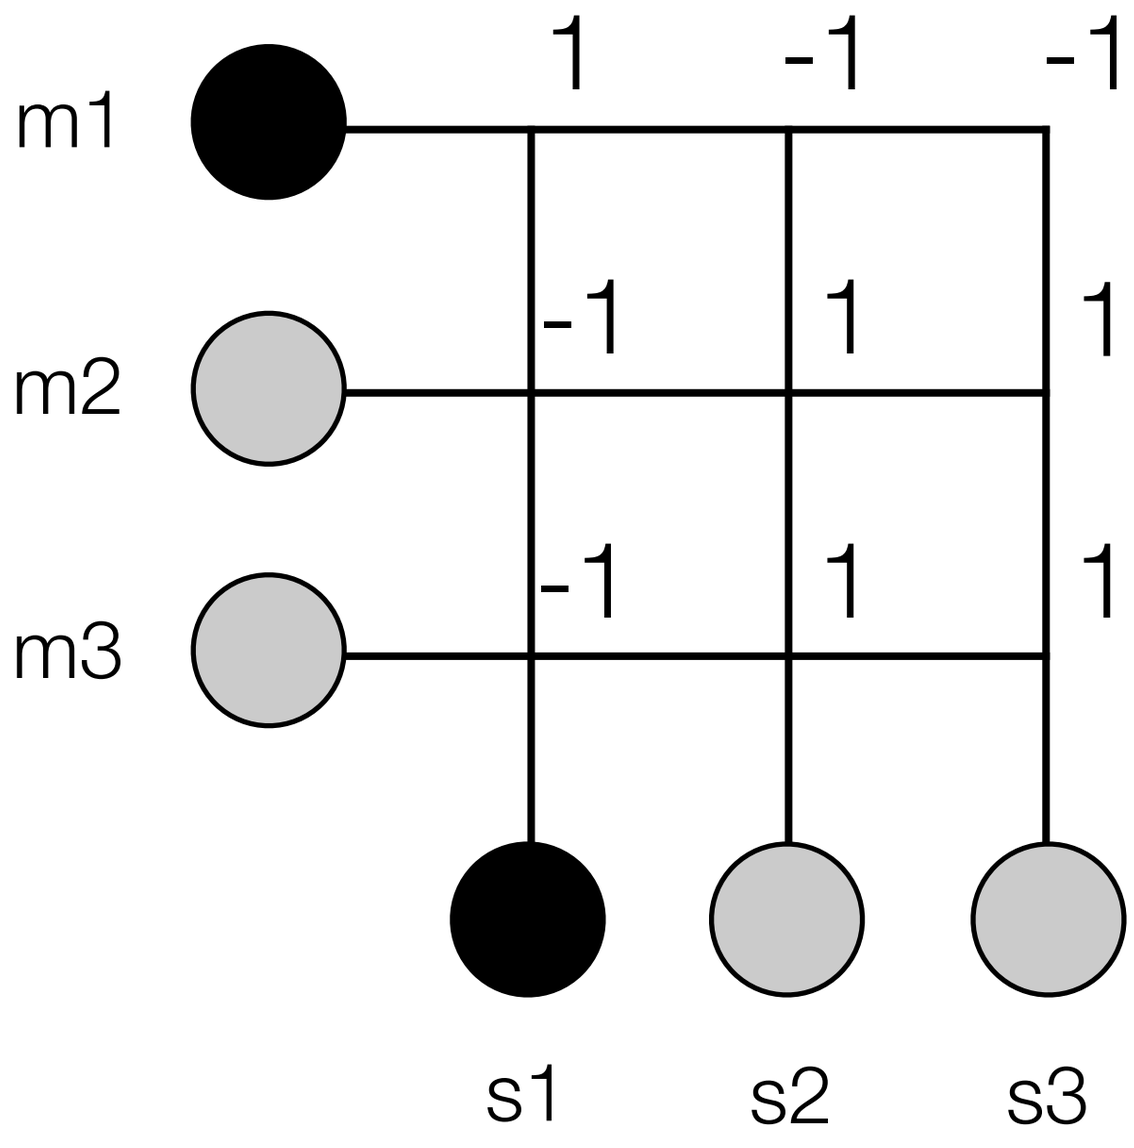

Working out bias

- A constructor rule: [+1, -1, -1, +1]

Production:

m2→?

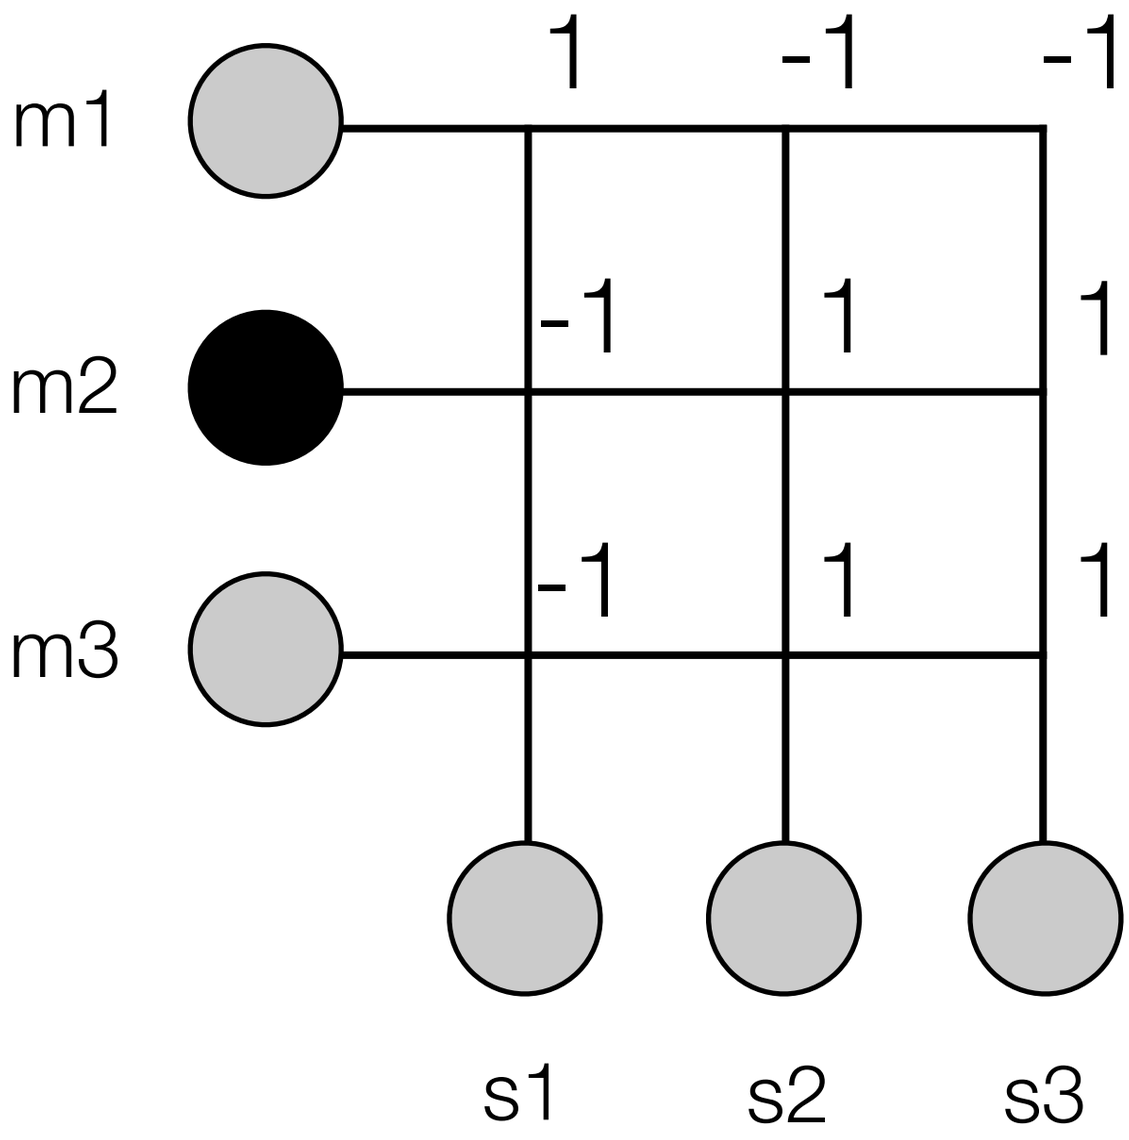

Working out bias

- A constructor rule: [+1, -1, -1, +1]

Production:

m2→s2 or s3

(*not* s1)

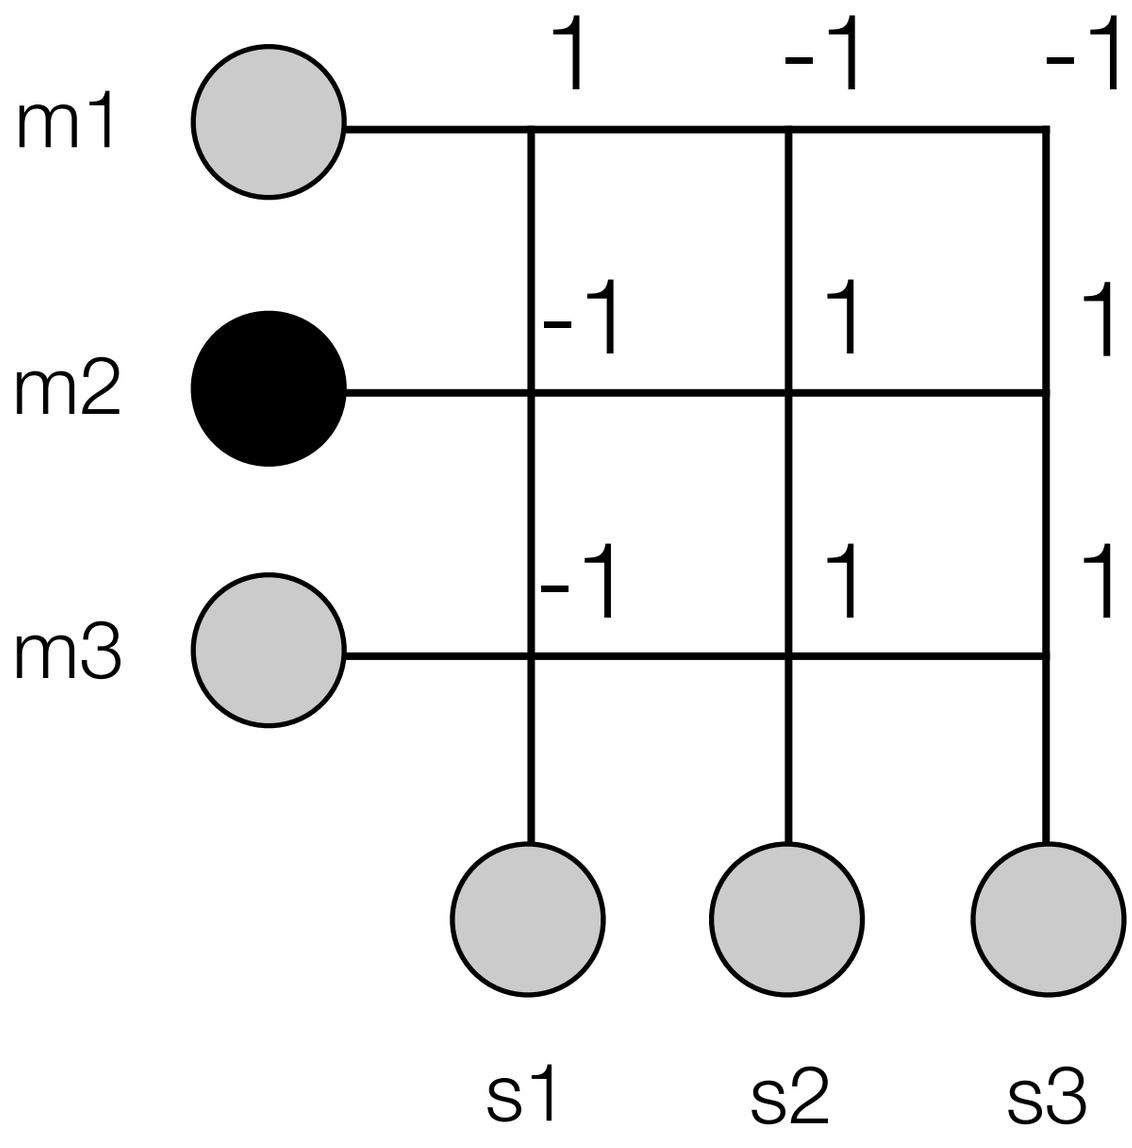

Working out bias

- Constructors in general: α > β & δ > γ
  After one exposure to m1→s1

Production:

m1→s1

m2→s2 or s3

m3→s2 or s3

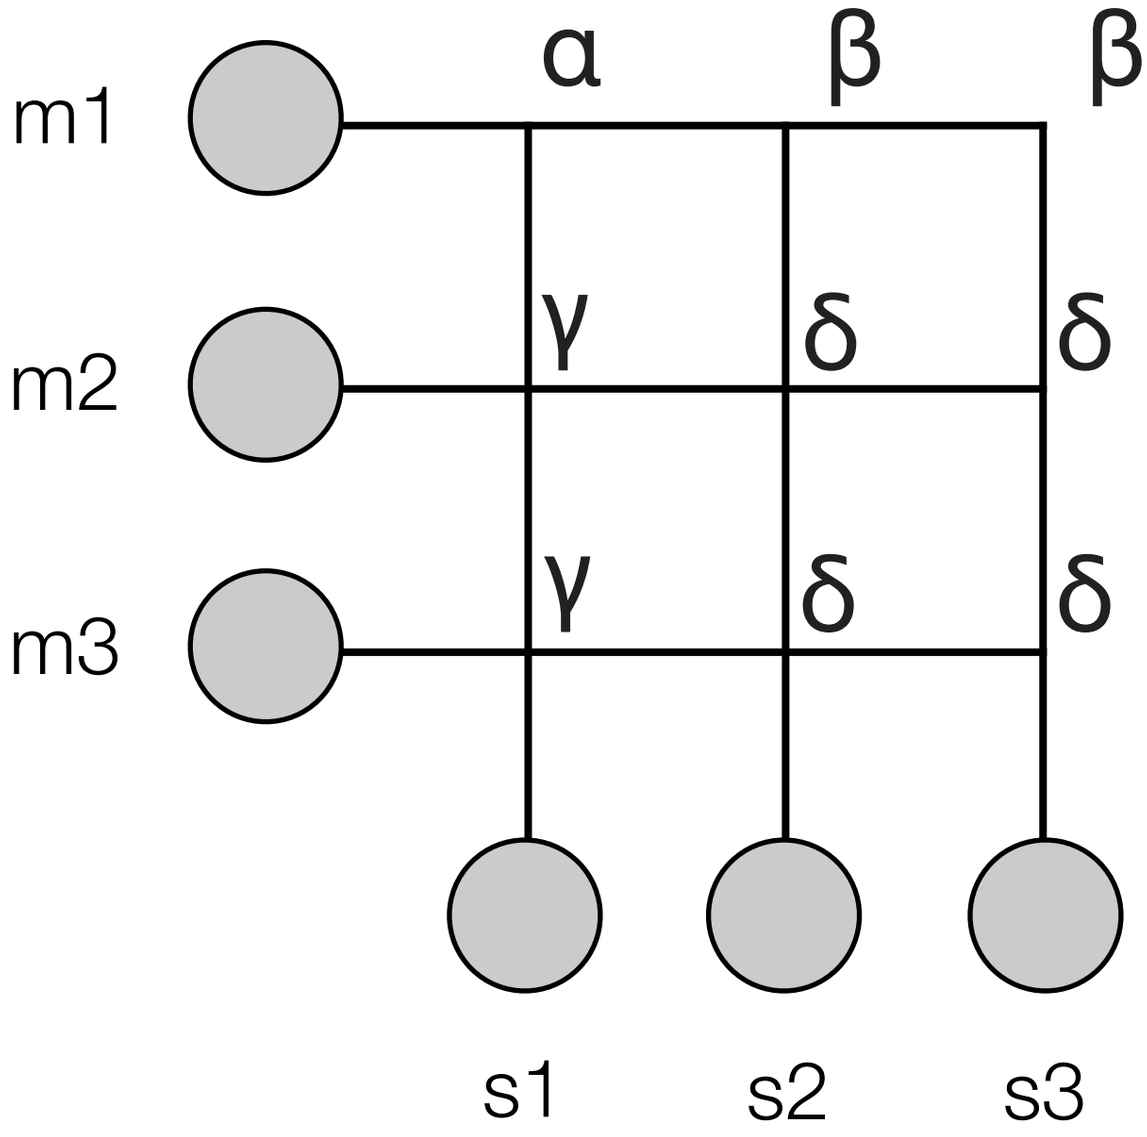

The constructor bias

- Constructors don’t like:

- One meaning to multiple signals

  because α > β
  bias against synonymy

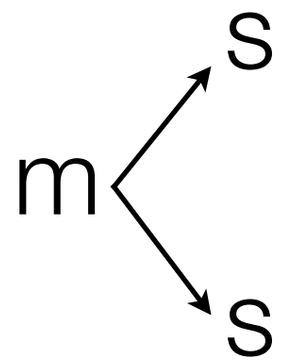

- Multiple meanings to one signal

  because δ > γ
  bias against homonymy

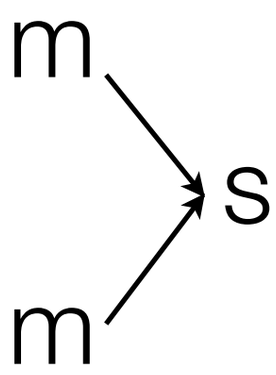

The constructor bias

- Constructors biased in favour of **one-to-one** mappings between meanings and signals

- Population's vocabulary changes over time to match this bias

- One-to-one systems happen to be optimal for communication

Working out bias

- A maintainer rule: [+1, 0, 0, 0]

Observation:
m1→s1

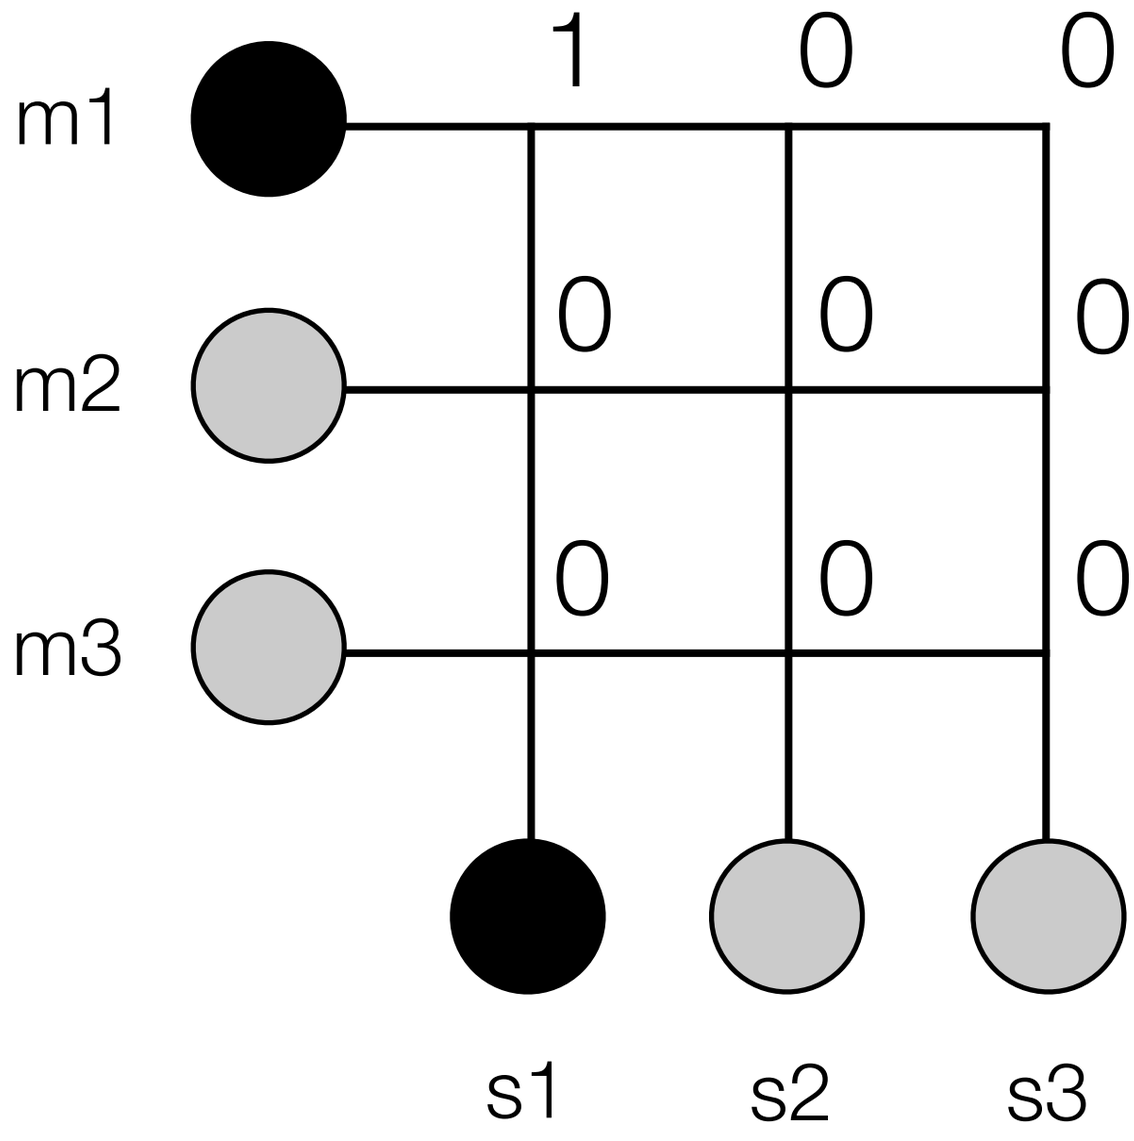

The maintainer bias

- Biased against synonymy

  because $\alpha > \beta$

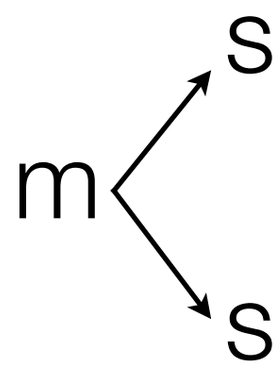

- **Neutral** with respect to homonymy

  because $\delta = \gamma$

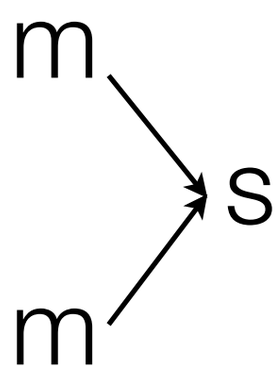

Rule: [1, 0, 0, 0]

Passes acquisition test? **Yes**

Maintenance: **Yes**

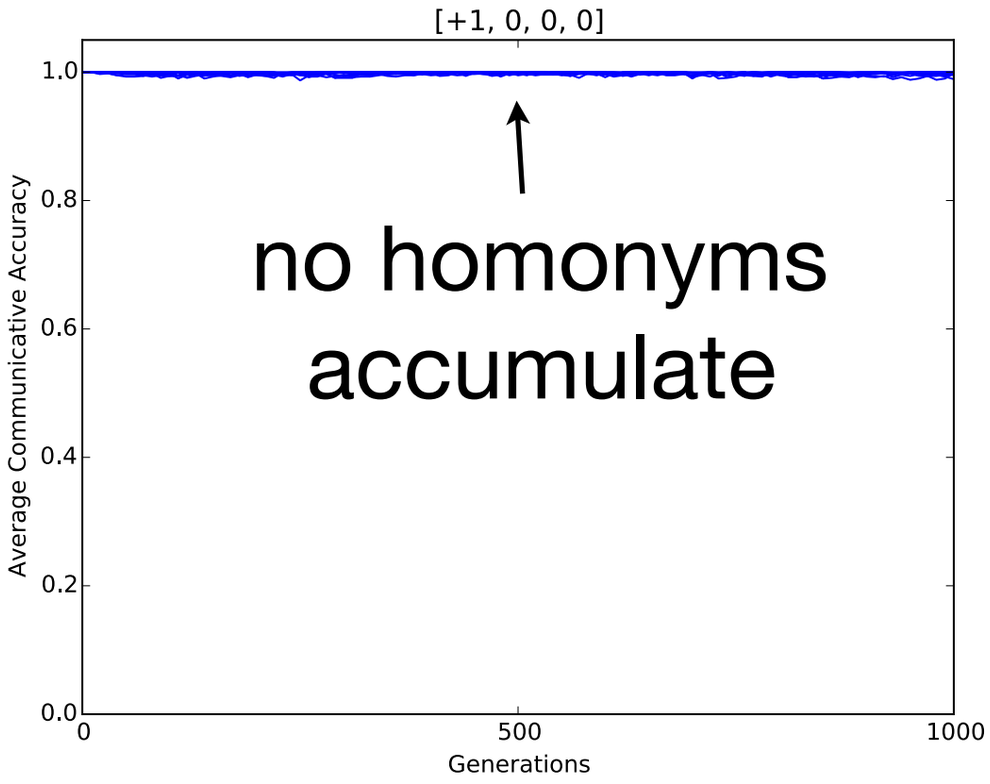

Construction: **No**

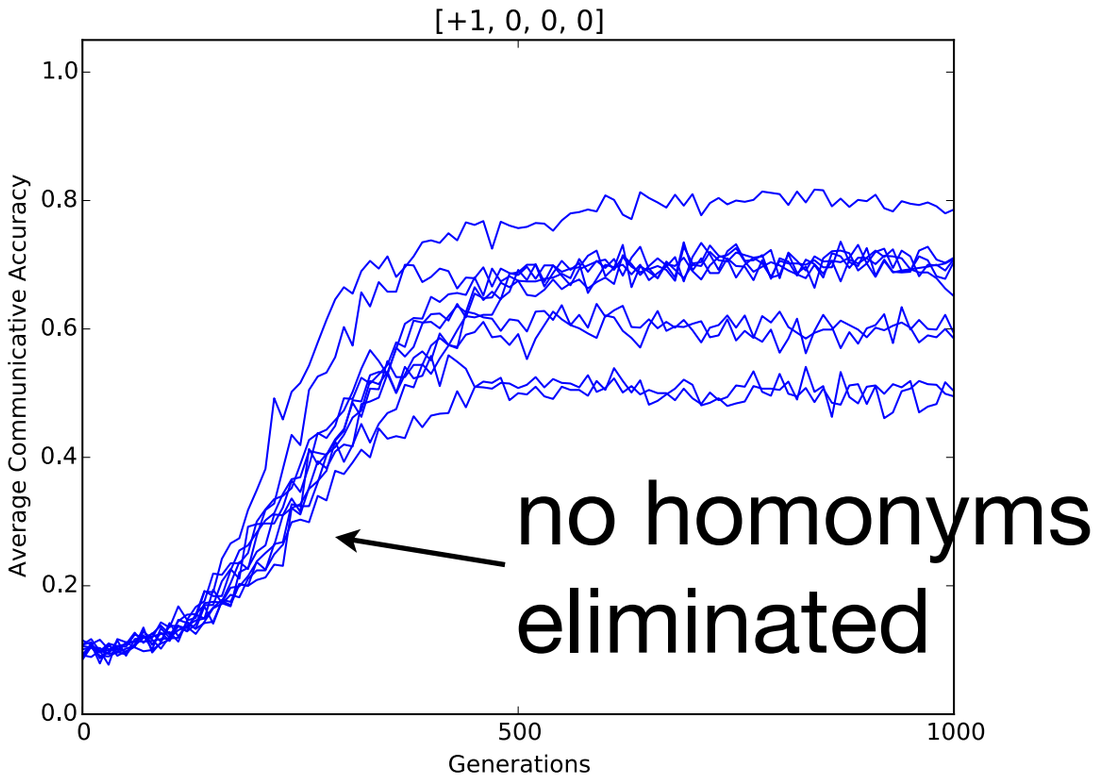

Working out bias

- A learner rule: [+1, -1, 1, 0]

Observation:
m1→s1

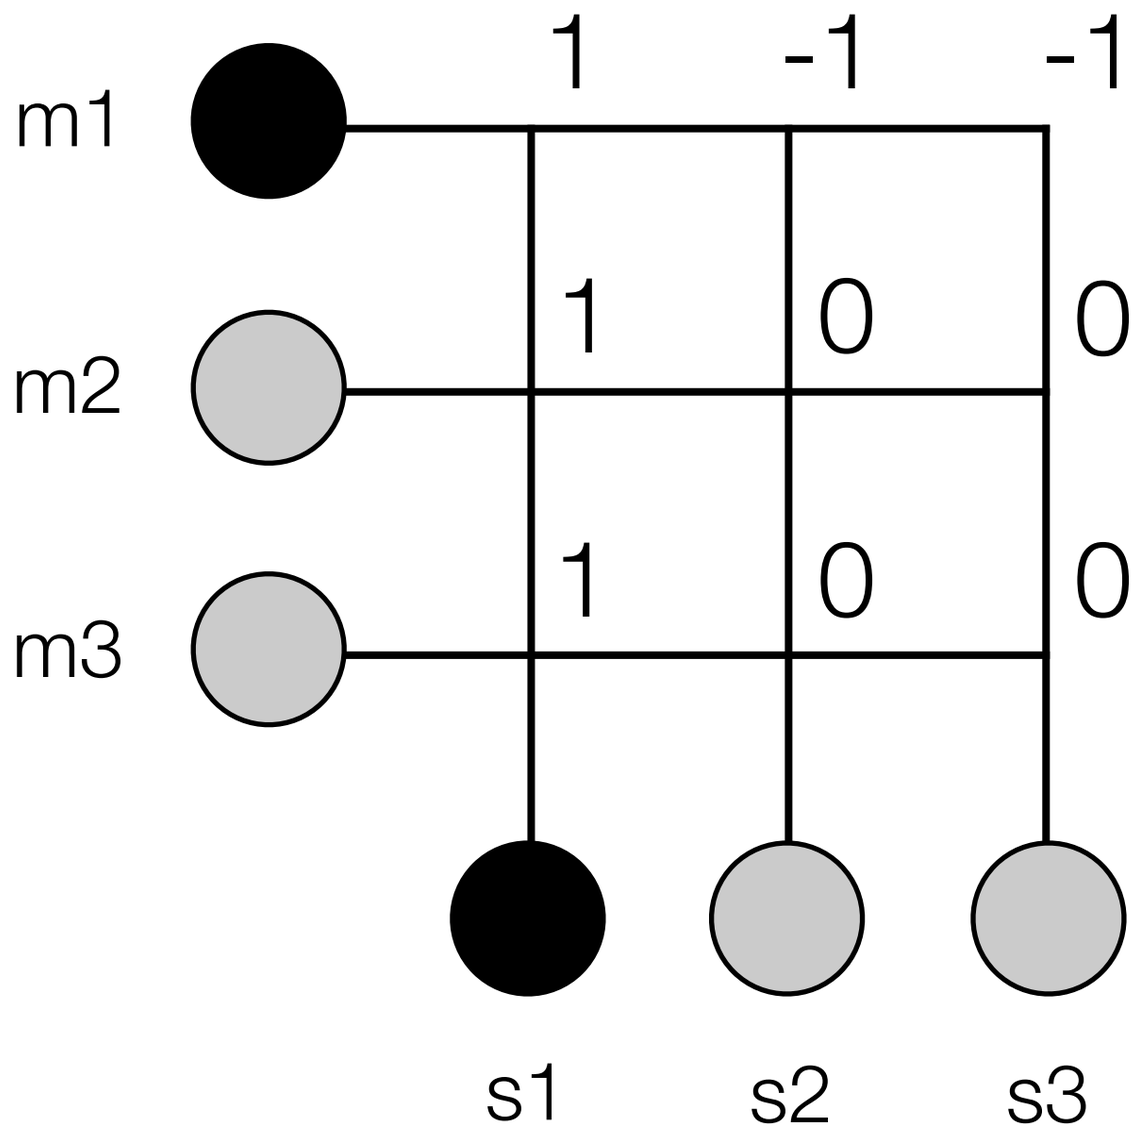

The learner bias (in most cases)

- Biased against synonymy

  because $\alpha > \beta$

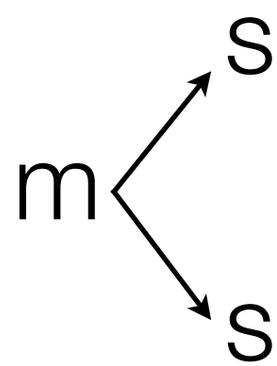

- Biased *in favour* of homonymy

  because $\delta < \gamma$

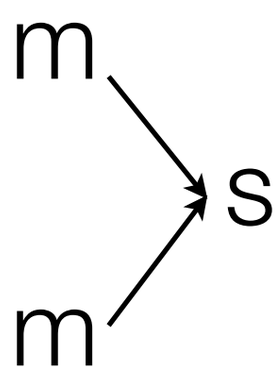

Rule: [1, -1, 1, 0]

Passes acquisition test? **Yes**

Maintenance: **No**

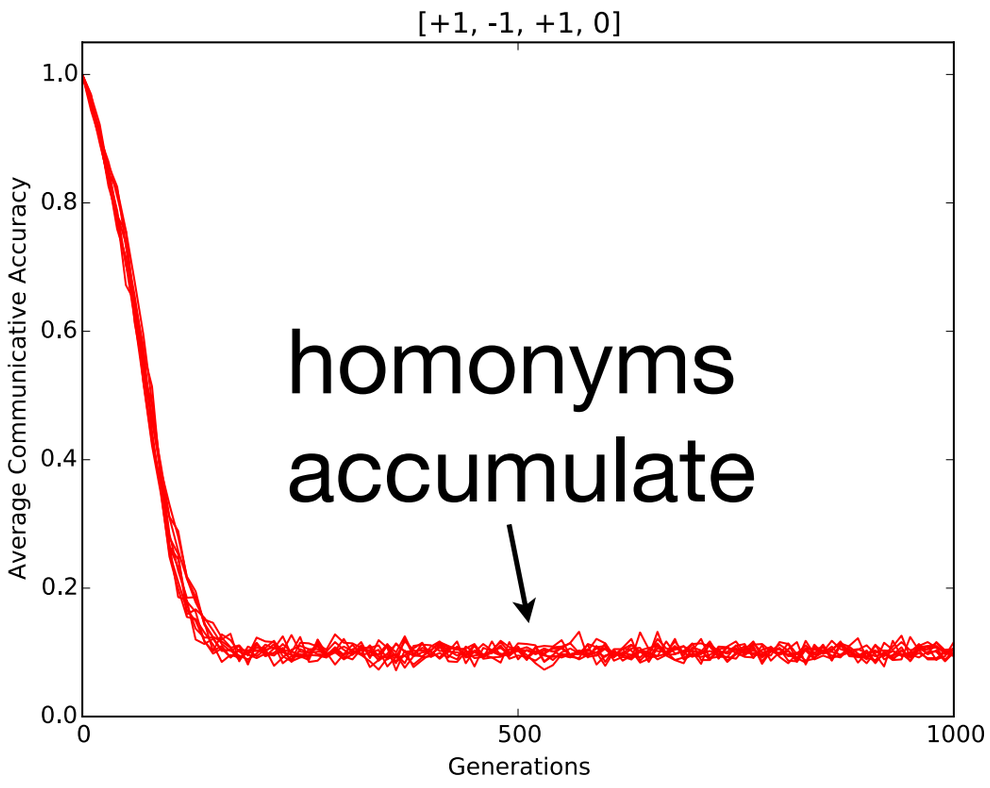

Construction: **No**

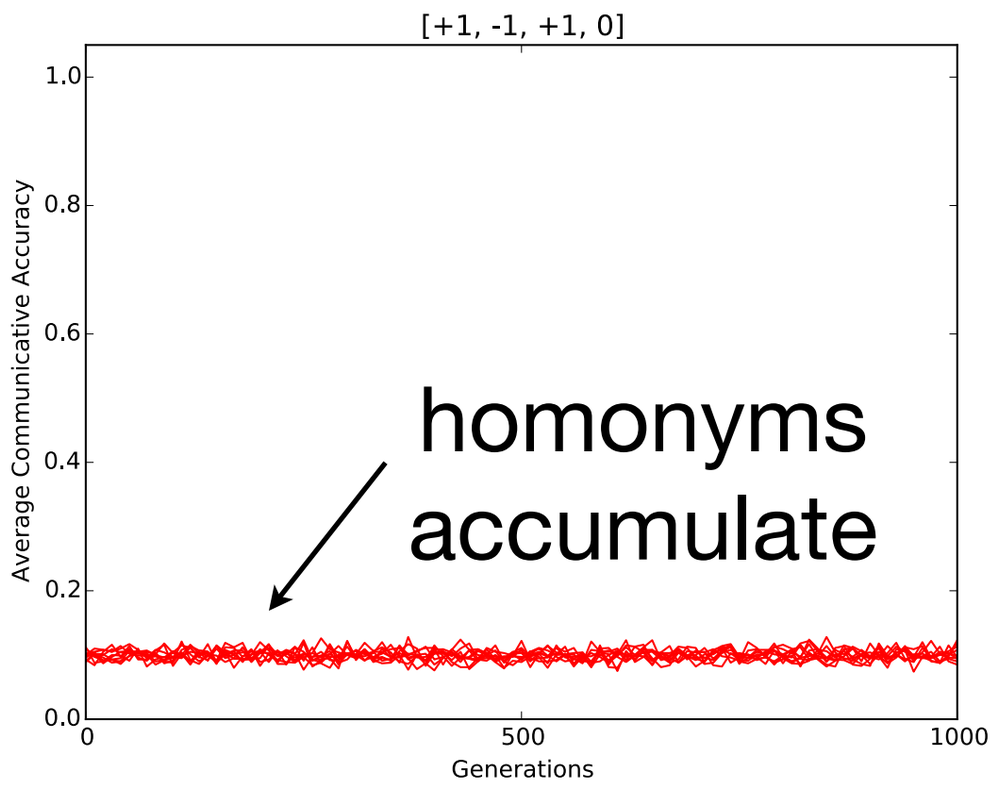

What about this rule? [0,-1,0,+1]

- A: it can neither maintain or construct
- B: it can maintain but not construct
- C: it can construct but not maintain
- D: it can maintain and construct

Rule: [0, -1, 0, +1]

Passes acquisition test? **Yes**

Maintenance: **Yes**

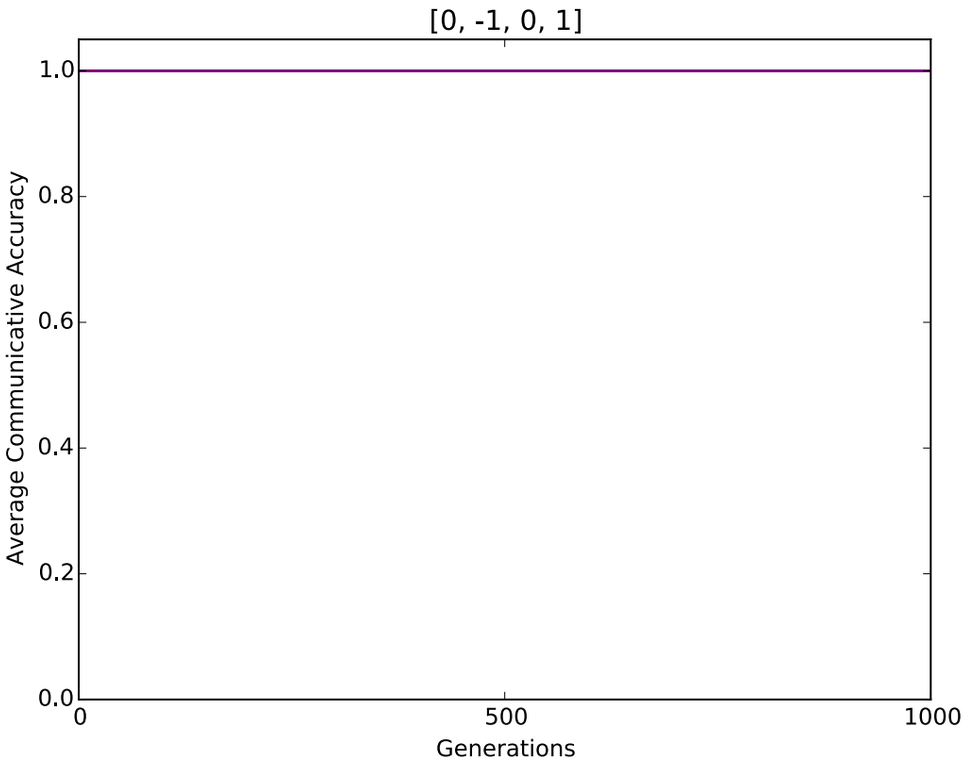

Construction: **Yes**

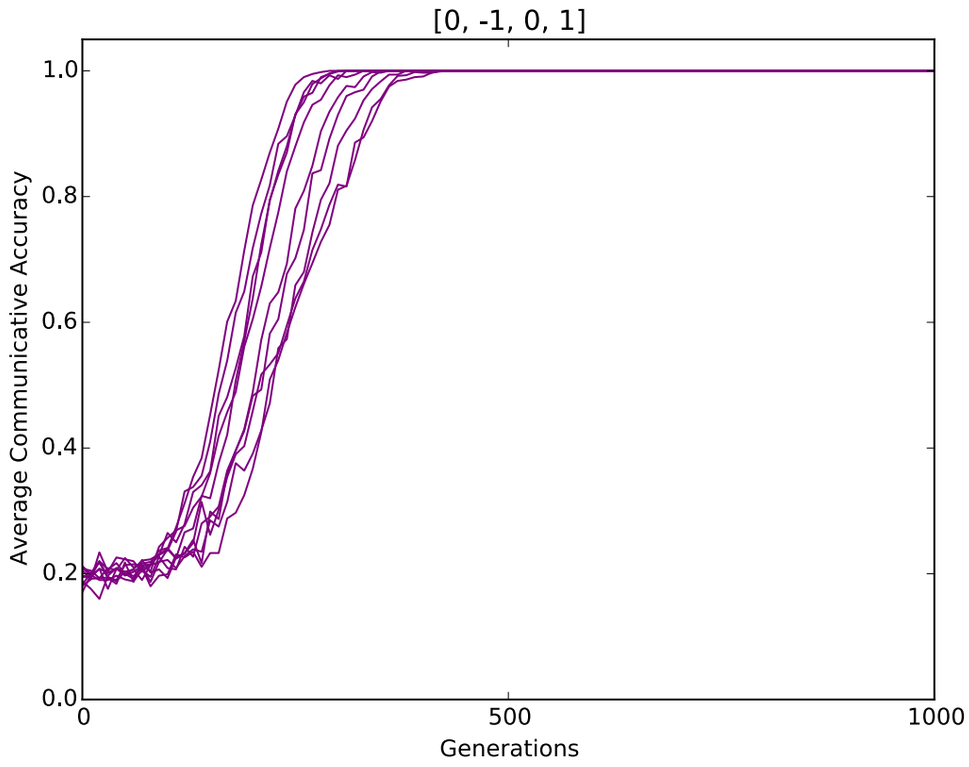

The constructor bias

- Constructors don’t like:

- One meaning to multiple signals

  because α > β
  bias against synonymy

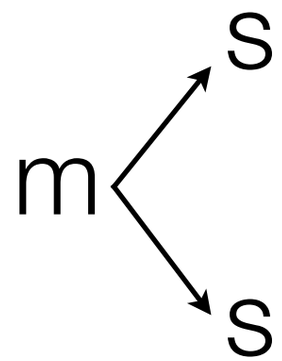

- Multiple meanings to one signal

  because δ > γ
  bias against homonymy

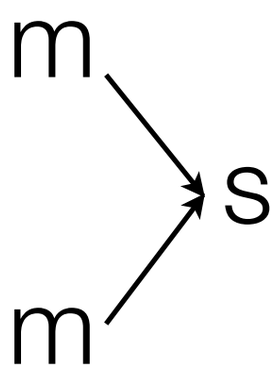

What about real humans?

- Experiment on children's learning bias
  Markman & Wachtel (1988) on synonymy

"Show me the fendle."

- Children pick the unfamiliar object given an unfamiliar word

Anti-synonymy bias (Mutual Exclusivity)

Before

banana

???

After (two possibilities)

banana

fendle

Homonymy bias (Doherty 2004)

- “... at the zoo, they saw a strange tapir from Brazil. Hamish thought the tapir’s long nose looked funny”

“Which one is the tapir in this story?”

Homonymy bias (Doherty 2004)

- "... at the zoo, they saw a strange **cake** from Brazil. Hamish thought the **cake's** long nose looked funny"

"Which one is the **cake** in this story?"

Anti-homonymy bias

Before

cake

???

After (two possibilities)

cak

c

cake

???

Children’s learning biases

- Children don’t like:
    - synonymy
    - homonymy
- They have the same biases as constructors in our simple model
- Populations of constructors evolve optimal communication systems
- Our model would predict that human vocabularies would be pushed in this direction **simply through iterated learning**, without additional functional pressures

Summary of the story so far, and what comes next

- Signalling systems (and languages) can evolve as a result of their transmission
    - We can model this
- The **biases** of learners shapes what evolves
- This potentially allows us to link findings about biases in learning at the individual level to predictions / observations about language at the population level
    - But caution (or better, a model) is required - the acquisition test here was misleading
- **Next up:** a class of models that allow us to be very clear and very precise about bias

A co-evolutionary hypothesis (Smith 2004)

Children's learning biases have evolved through natural selection, because they're good for communication.

- Examine this idea using our model
- Two central assumptions:
  - Weight update rule is given by a genotype
  - Better communicators breed more

Invasion of the mutants

- Smith (2004) plays **constructors**, **maintainers**, and **learners** off against each other

- Create a population mainly made up of one type, but with a small number of another type (the mutant)

- Agents inherit both the communication system (by cultural transmission), and their learning strategy (by genetic transmission)

- Both culture and biology evolve

- If selection is based on communicative success, which mutants will invade?

Surprising result: evolution is hard

- Constructors don't often invade, *even though it would increase the fitness of the population if they did*

- Two problems:
  - Need a lot of mutants before they start to have a good effect on the population's language...
  - ...and even then, there's a time-delay before the good language evolves culturally.

- Speculative conclusion: human learning biases *haven't* evolved only for communication.

Summary

- Smith (2002, 2004) look in detail at how learning bias can give us (or fail to give us) language

- Brings together 3 complex processes in one model:

  Learning
  Cultural transmission
  Biological evolution

- Highlights the crucial importance of the second of these three