Exploration and Exploitation

PJ Lamberson
University of Michigan

James G. March

Google Scholar BETA

Exploration and **Exploitation** in Organizational Learning
JG March - Organization Science, 1991 - JSTOR
Exploration and **Exploitation** in Organizational Learning. James G. March. **...** 82. **EXPLORATION** & **EXPLOITATION** IN ORGANIZATIONAL LEARNING The pattern is clear. **...**
Cited by 2381 - Related Articles - Web Search

Exploration

"... includes things captured by terms such as search, variation, risk taking, experimentation, play, flexibility, discovery, innovation."

Exploitation

"... includes such things as refinement, choice, production, efficiency, selection, implementation, execution."

The Tradeoff

All Exploration: "... the costs of experimentation without any of its benefits." Undeveloped ideas, little distinctive competence.

All Exploitation: Locked-in to suboptimal equilibria (local maxima). Can't adapt to changing circumstances.

The Tradeoff

Exploration: long-term, risky, uncertain.

Exploitation: short-term, immediate, certain benefits.

The Tradeoff

An organization, such as a firm, a government or a political party has to choose how much of their resources to allocate to each of these activities.

The Tradeoff

An organization, such as a firm, a government or a political party has to choose how much of their resources to allocate to each of these activities.

Explicit and Implicit Choices

Explicit: Calculated decisions about investment and strategy.

Implicit: Organizational forms and customs. Procedures for accumulating and reducing slack, search rules and practices, the ways in which targets are set and changed, incentive systems, hiring practices ...

Multiple Levels

- Individual level
- Organizational level
- Social system level

“Understanding the choices and improving the balance between exploration and exploitation are complicated by the fact that returns from the two options vary not only with respect to their expected values, but also with respect to their variability, their timing, and their distribution within and beyond the organization. Processes for allocating resources between them, therefore, embody intertemporal, interinstitutional, and interpersonal comparisons, as well as risk preferences.”

Feedback to exploitation occurs much more quickly. Increasing returns can lead to lock-in at a sub-optimal equilibria.

"... these tendencies to increase exploitation and reduce exploration make adaptive processes potentially self-destructive."

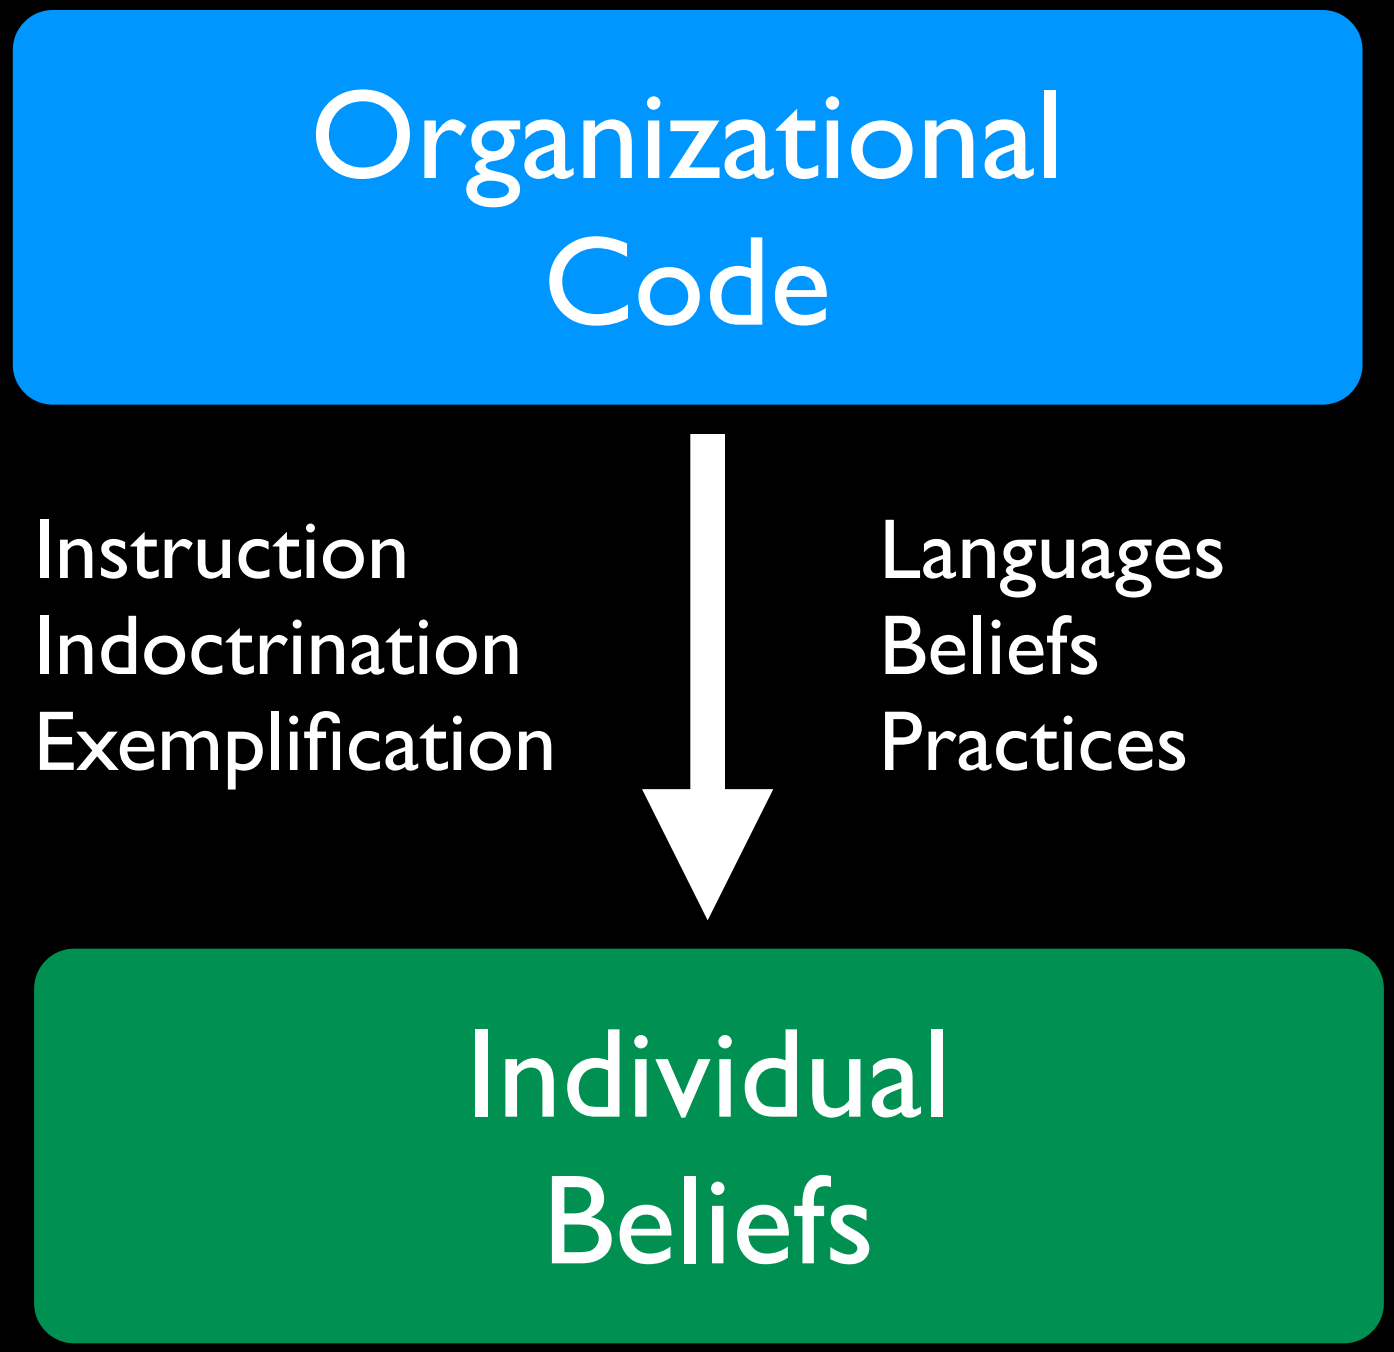
Organizational Code
Instruction
Indoctrination
Exemplification
Languages
Beliefs
Practices
Individual Beliefs

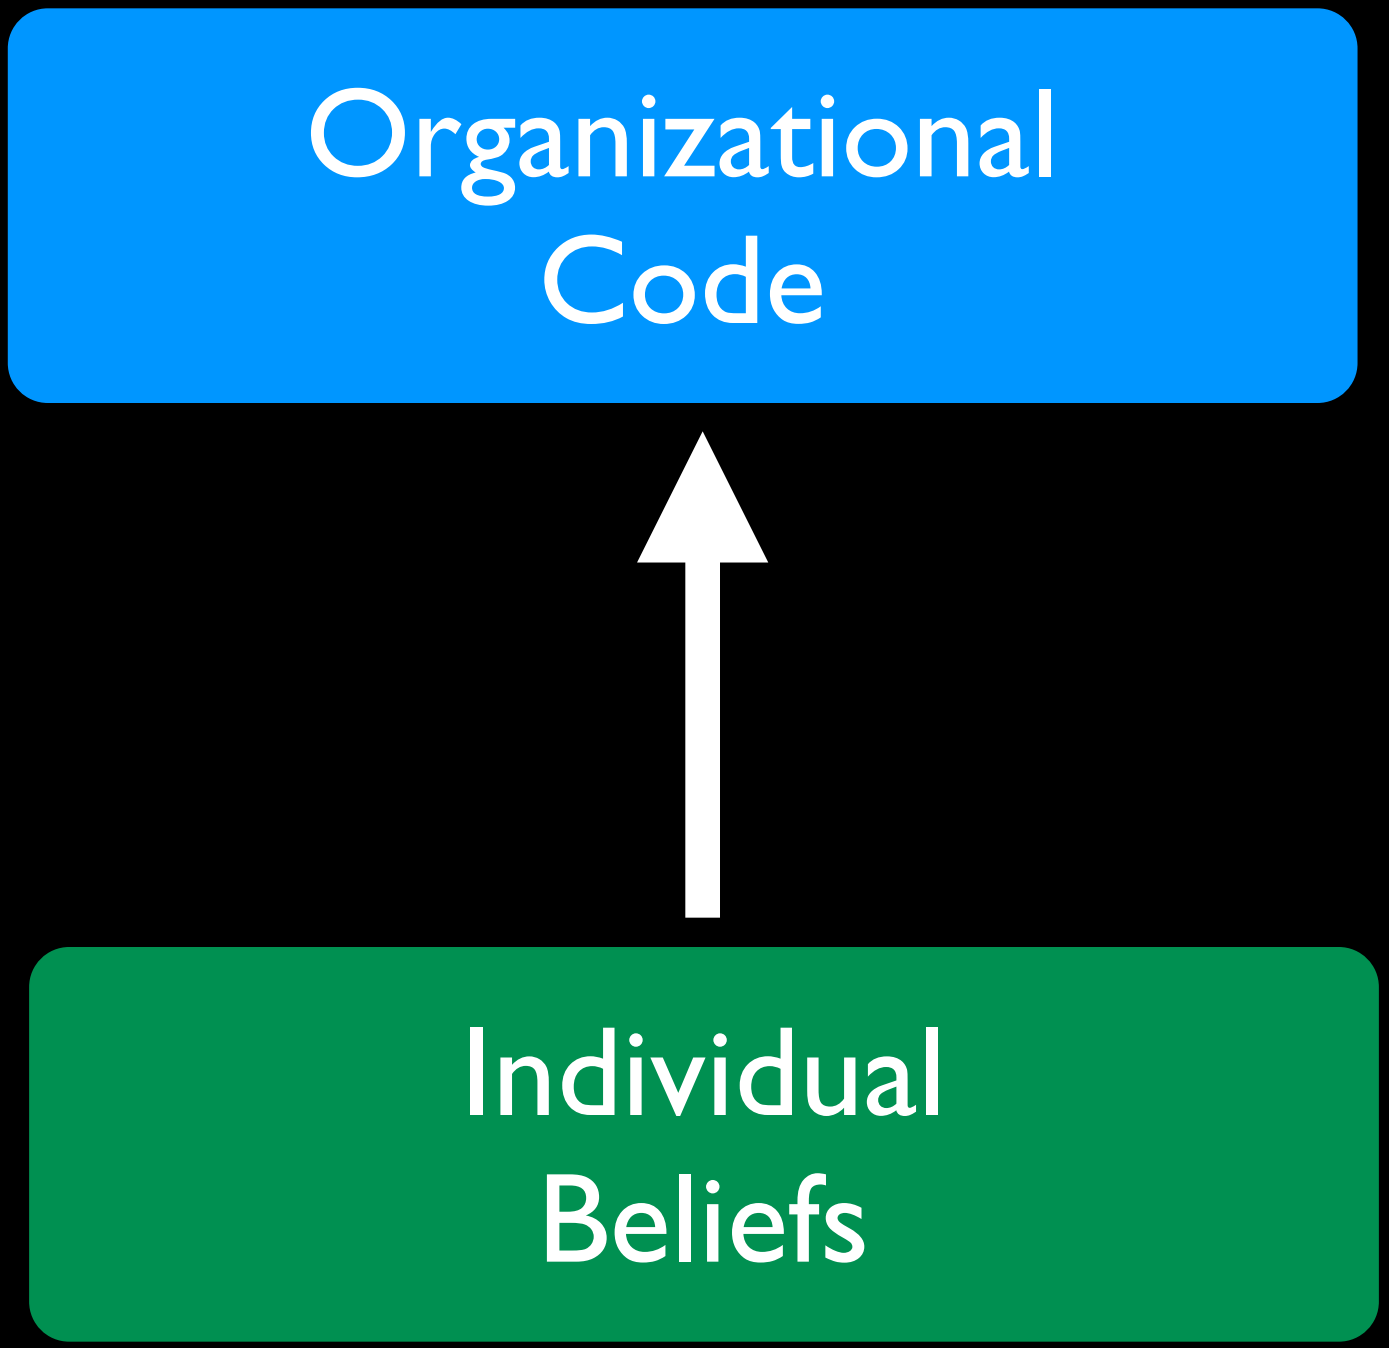
Organizational Code
Individual Beliefs

The Model

reality

$$\text{Reality}=(\underbrace{1,1,-1,1,\ldots,1}_{m})$$

Probability of any given dimension being 1 is .5

The Model

the code and beliefs

$$\text{Organizational code}=(\underbrace{1,0,-1,0,\ldots,1}_{m})$$

$$\text{Individual beliefs}=(\underbrace{1,0,-1,0,\ldots,1}_{m})\Big\}\,n$$

The Model

socialization

In each period, for each dimension that the code is non-zero, with probability p_1 an individual updates her beliefs to match those of the code.

p_1 is the socialization parameter.

The Model

learning by the code

The code adapts to the beliefs of individuals whose beliefs correspond with reality on more dimensions than the code does.

This part gets confusing!

The Model

learning by the code

First, assign everyone a score, which is given by the number of dimensions that their beliefs agree with those of the code.

Second, disregard everyone that has a lower score than that of the code. Call everyone that is left, the *superior group*.

The Model

learning by the code

Make a new set of beliefs, called the *superior majority beliefs* given by taking the view of the majority of the people in the superior group on each dimension.

For each dimension, if the code agrees with the superior majority beliefs, don't change it.

The Model

learning by the code

If the code disagrees with the superior majority beliefs, change the code to those of the superior majority with probability $1-(1-p_2)^k$, where k is the number of people in the superior group that share the superior majority belief minus the number who do not.

The Model

initialization

Reality is randomly 1 or -1 on each dimension.

The code is all 0s.

Individual beliefs are randomly 1, 0 or -1 on each dimension.

The Model

initialization

So, initially the code is neutral, and on average the individuals no nothing.

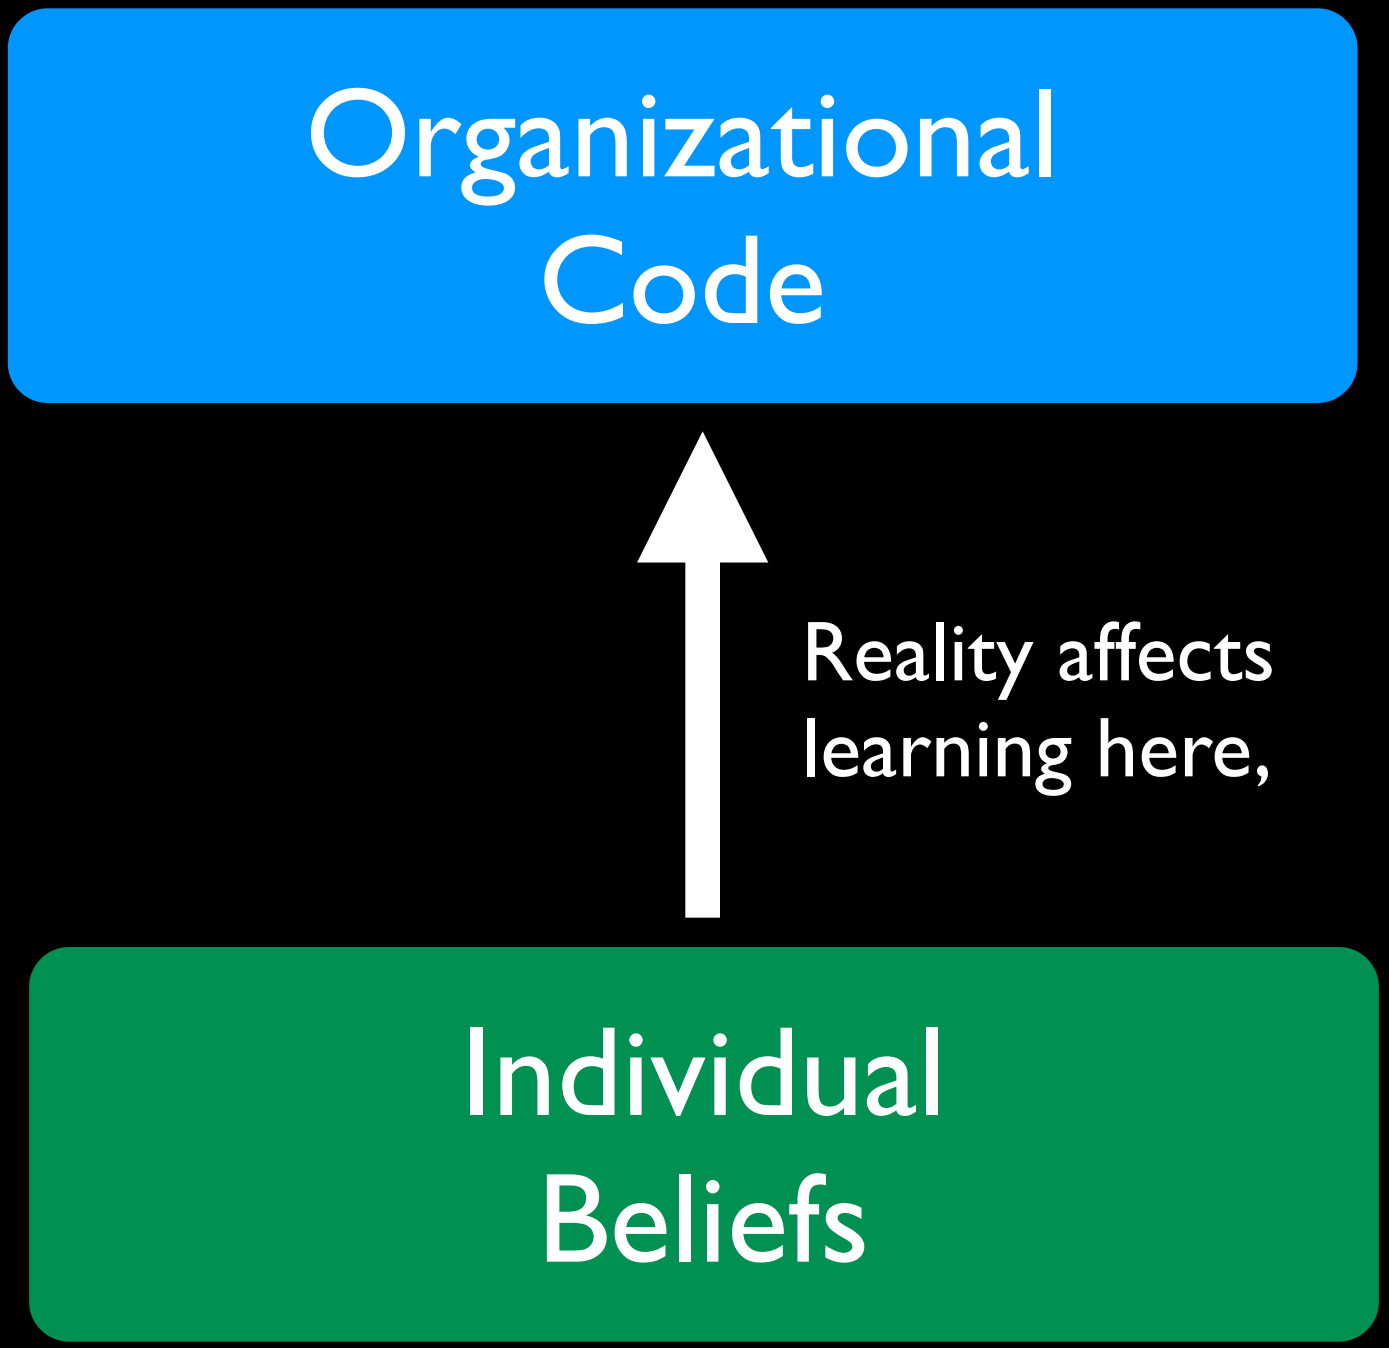
Organizational Code
Reality affects learning here,
Individual Beliefs

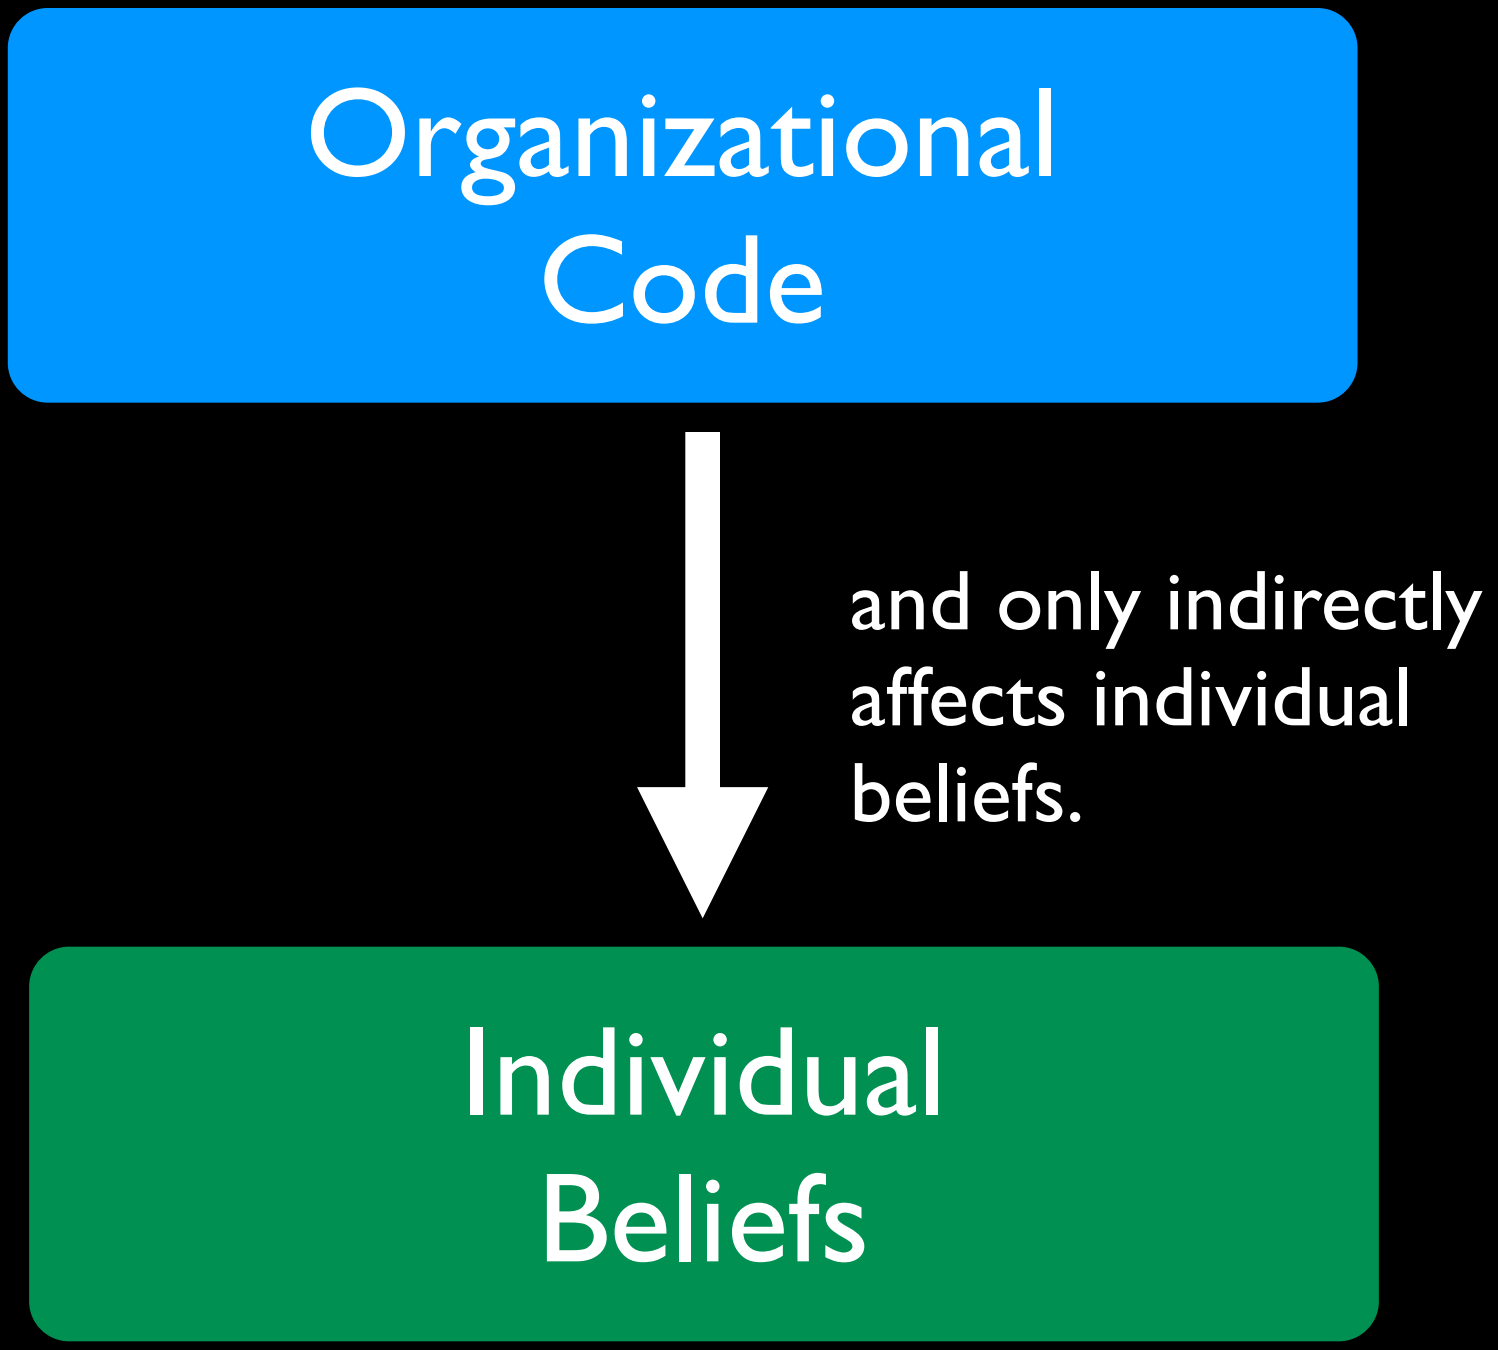
Organizational Code
and only indirectly affects individual beliefs.
Individual Beliefs

Neither the organization, nor the individuals possess any knowledge of reality. The organization knows what individuals are closer to reality than it (but not on which dimensions!). That's it! Both individuals and the code can copy false beliefs from one another.

All changes lead individuals and the code to be more similar. Once the code and all individuals share the same belief, nothing can change.

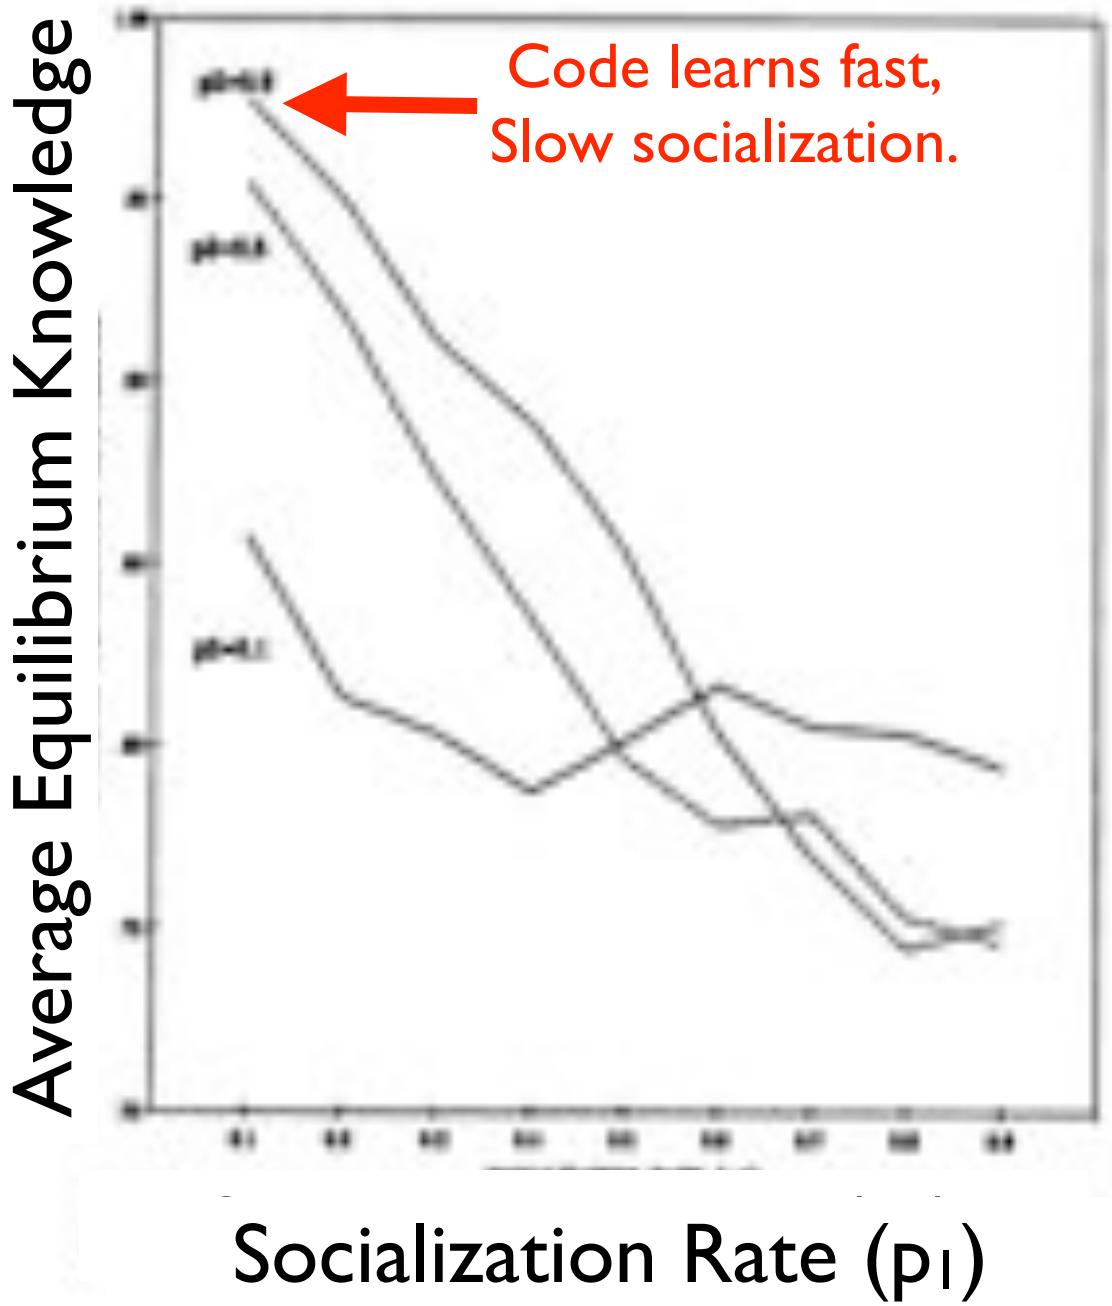

Slower socialization is better.

When socialization is slow, faster code learning is better.

The best is when the code learns fast, but socialization is slow.

“The gains to individuals from adapting rapidly to the code (which is consistently closer to reality than the average individual) are offset by second-order losses stemming from the fact that code can learn only from individuals who deviate from it. Slow learning on the art of the individuals maintains diversity longer, thereby providing the exploration that allows the knowledge found in the organizational code to improve.”

Slow Learners and Fast Learners

So far, everyone learns from the code at the same rate (p_l). Now, suppose we have *fast learners* (p_l =.9) and *slow learners* (p_l =.1).

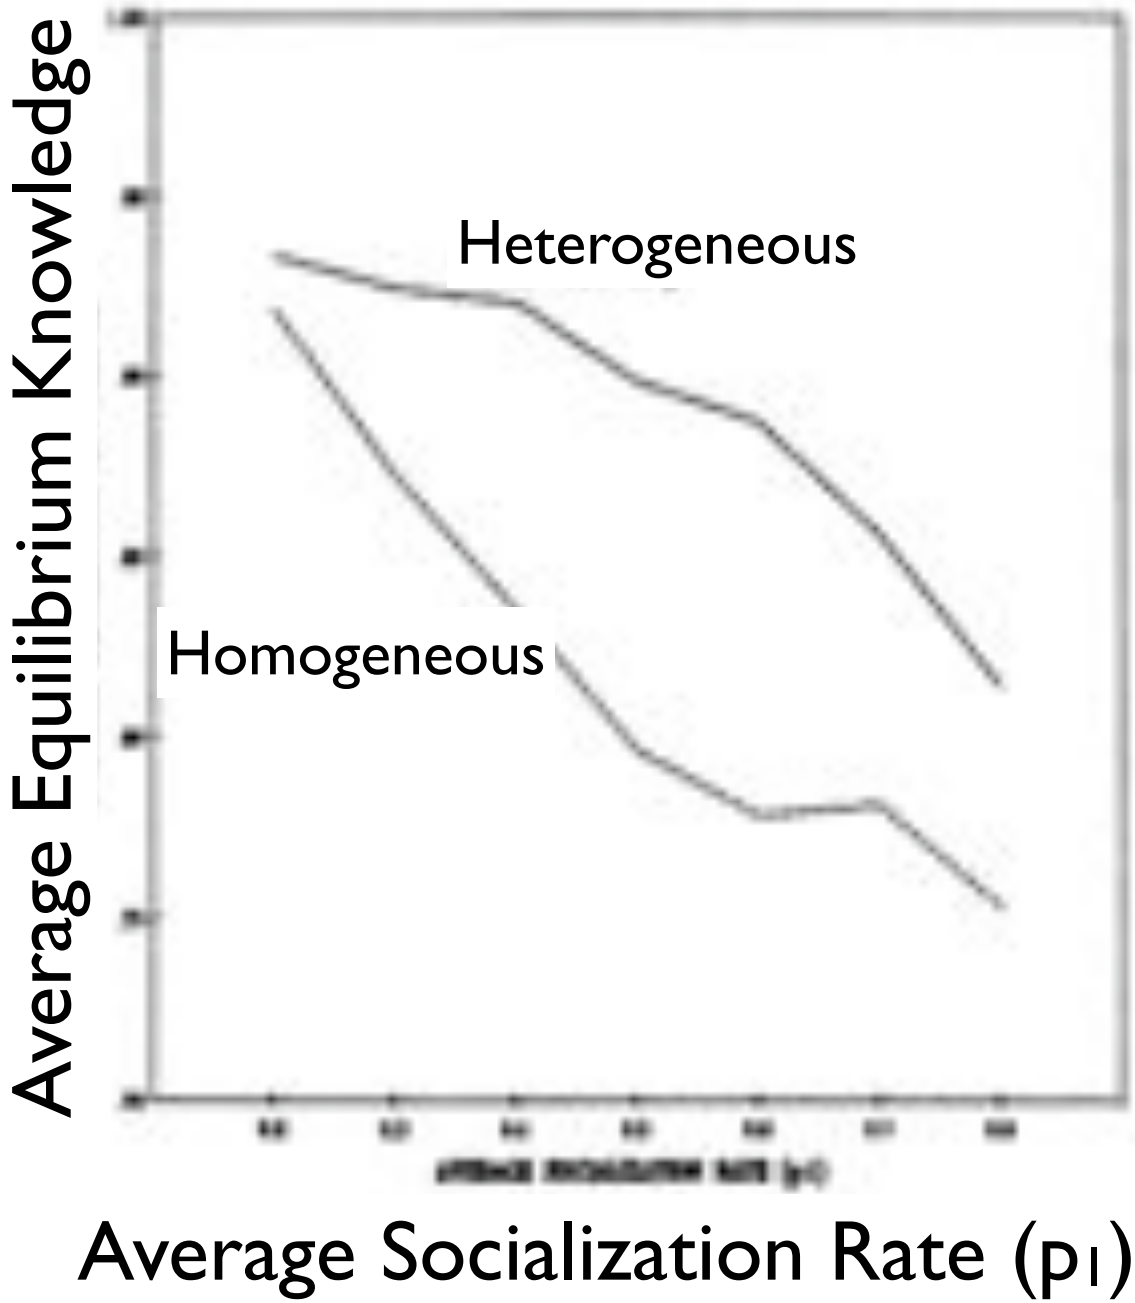

Slower socialization is better.

Given an average socialization level, it is better to have heterogeneous individuals than homogeneous.

The code gains most of its knowledge from the slow learners, and that knowledge is the realized by the fast learners.

The fast learners benefit a lot from the slow learners, but the slow learners do not benefit from the fast learners.

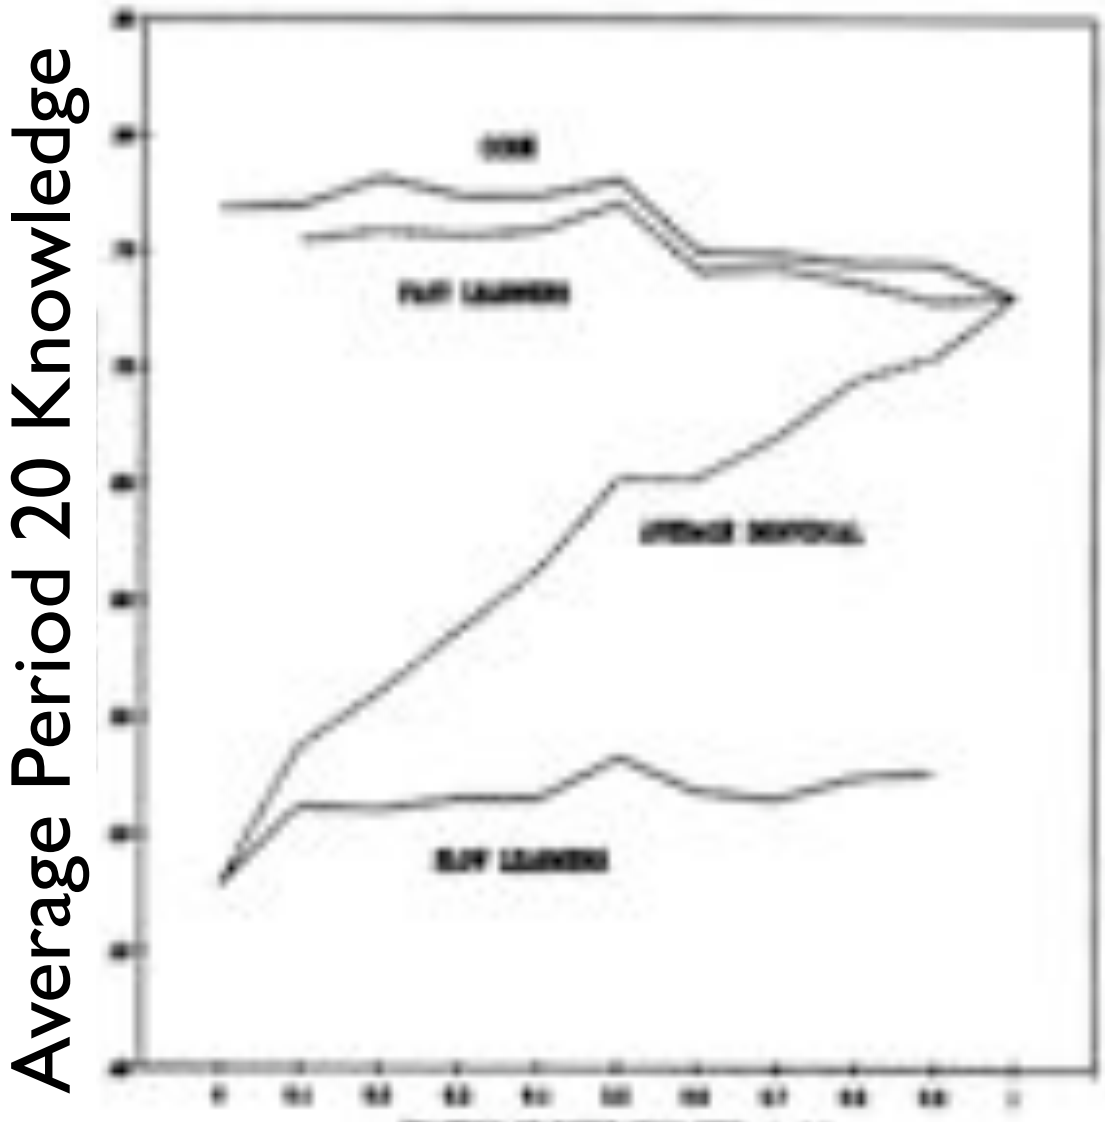

Who would want to be a slow learner?

You're fired!

turnover

Suppose that each time period, each individual has a probability p3 of leaving and being replaced by a new employee with randomly assigned individual beliefs.

This should decrease average individual knowledge, because on average individuals in the organization are more knowledgeable than random people.

How does it affect the code?

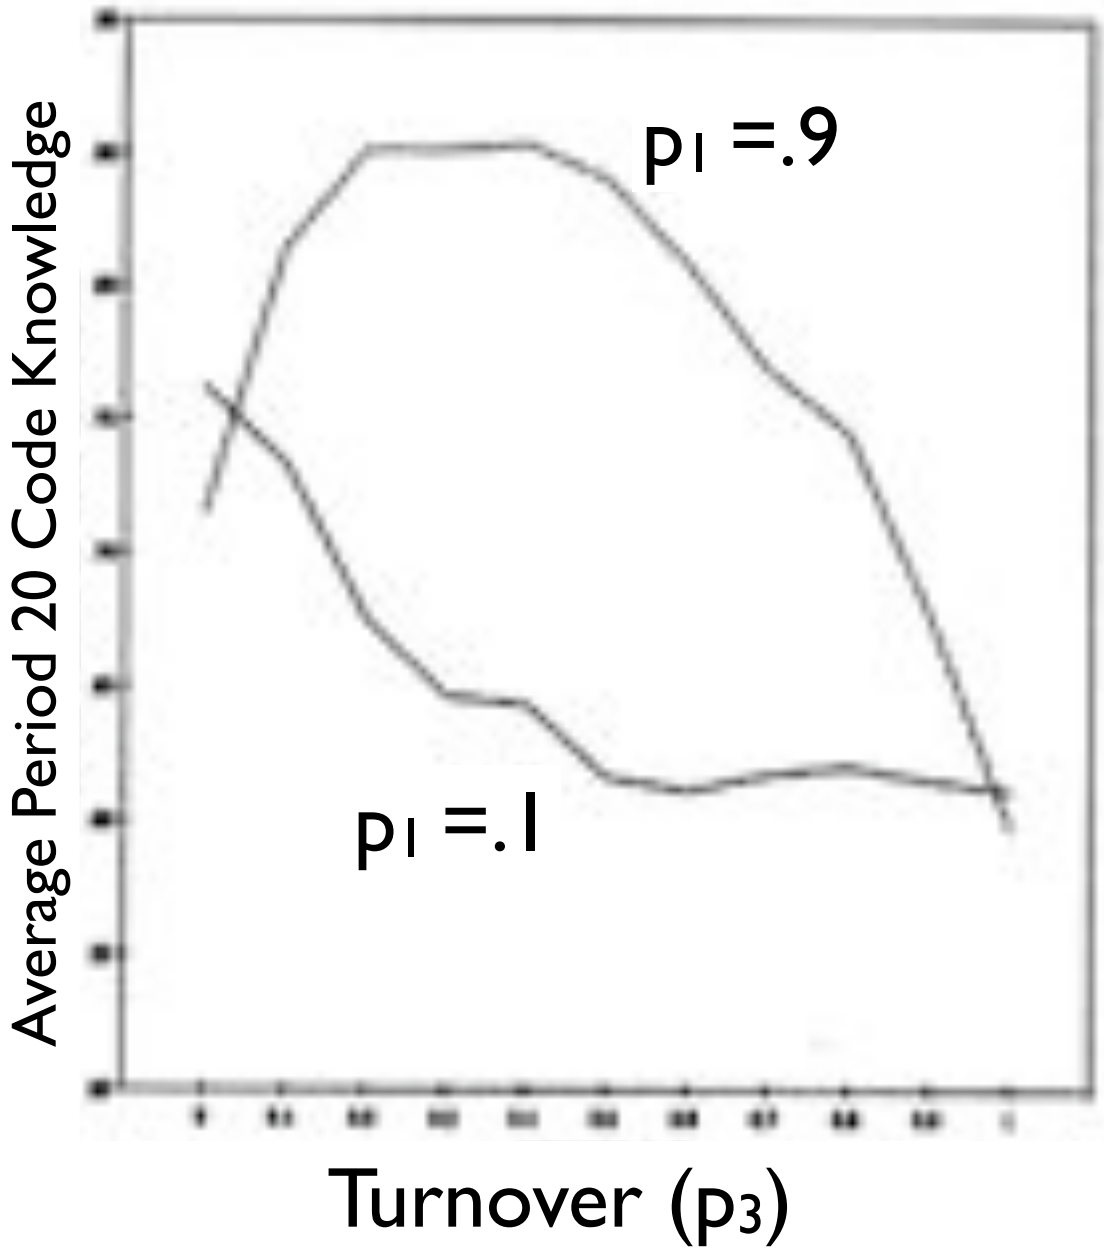

If socialization is slow, turnover is bad.

If socialization is fast, some turnover is good, but too much is bad again.

“Old-timers, on average know more, but what they know is redundant with knowledge already reflected in the code ... Novices know less on average, but what they know is less redundant with the code and occasionally better, thus more likely to contribute to improving the code.”

Turbulence

The value of any given dimension shifts randomly in any given time period with probability p_4.

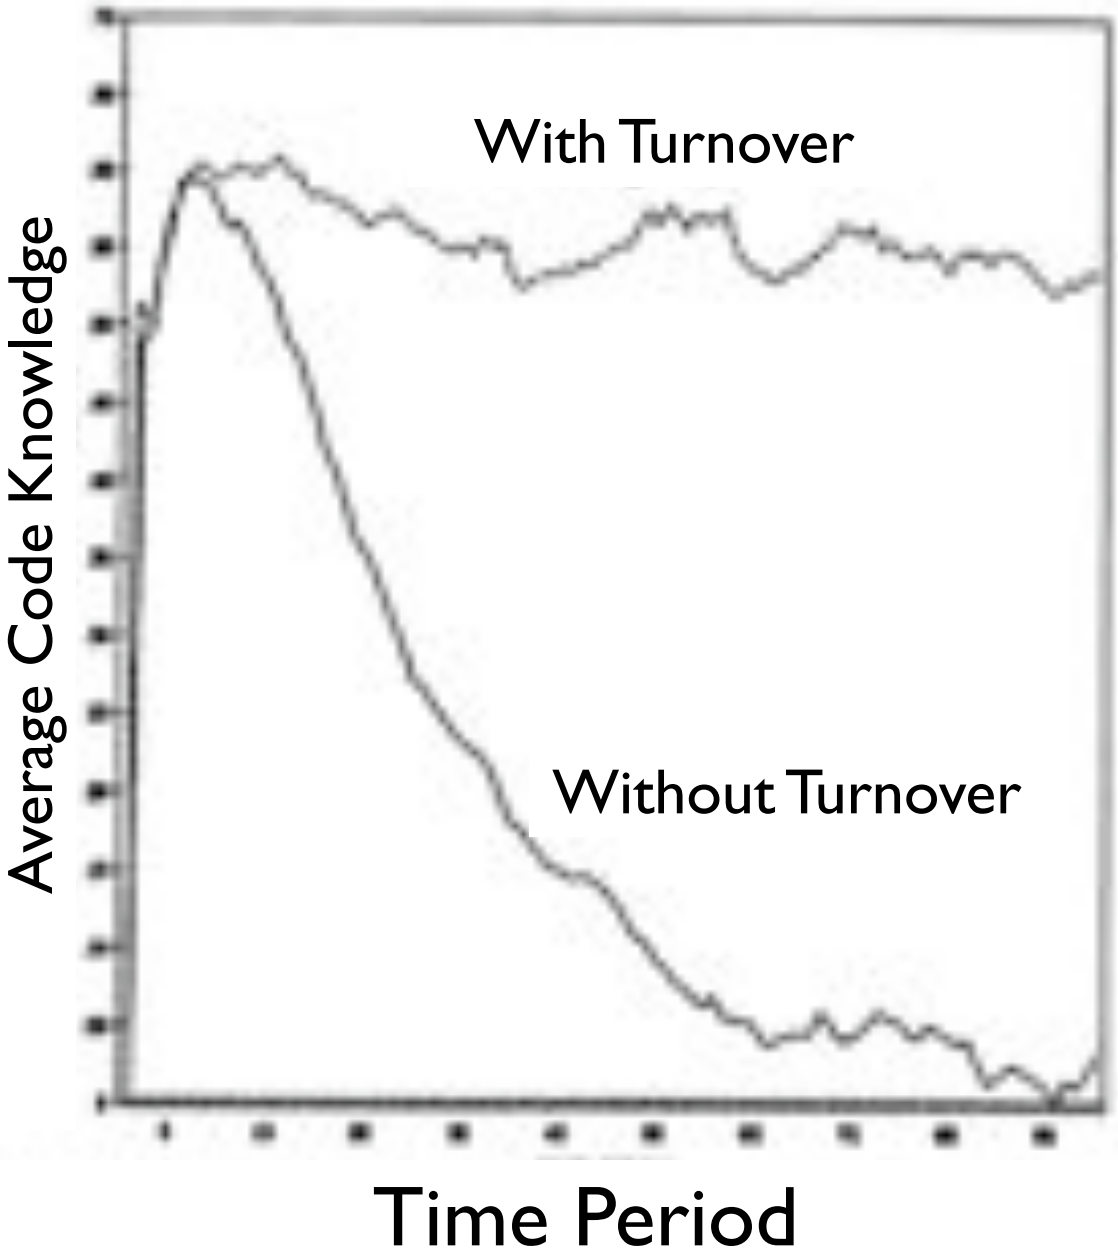

Organizations without turnover are locked-in and cannot adapt to changes in reality.

So, why has this paper been cited 2381 times?

Good Modeling Practices

- Keep the model simple
- Focus on the science, not the computer
- The old computer test
- Avoid black boxes

- Miller and Page, *Complex Adaptive Systems*

Good Modeling Practices

- Nest your models
- Have tunable dials
- Reward the right things

- Miller and Page, *Complex Adaptive Systems*