Rainer Hegselmann,
Frankfurt School of Finance & Management,
Department of Computational Science & Philosophy

Polarisation in the bounded confidence model: A revised view

Acknowledgement
This presentation uses insights and results from long lasting collaborations with Igor Douven (CNRS/Sorbonne), Ulrich Krause (Bremen), Sascha Kurz (Bayreuth), and Jörg Rambau (Bayreuth).

Workshop *Measuring, Modeling and Mitigating Opinion Polarization and Political Cleavage*
ETH Zürich, September 13-15, 2023

Structure of the talk

1. The bounded confidence model: Idea and short analysis
2. A new analytical approach: Given a certain start distribution X(0), how many *different* BC-processes $\langle X(0), \epsilon \rangle$do exist?
3. ϵ-Switches of equidistant start distributions with $n = 2, 3, \ldots, 50$ agents
4. ϵ-Switch diagrams
5. Lessons and perspectives

Disclaimer: *No empirically calibrated model!*

Instead: *Computer aided* thinking, theorising, speculating, thought experiments on…

- … *possible* mechanisms, effects, processes, and *possible* components of *possible* partial explanations,
- … the *possible* interplay of mechanisms,
- … *possible* interventions.

with the help of a creation machine for artificial worlds

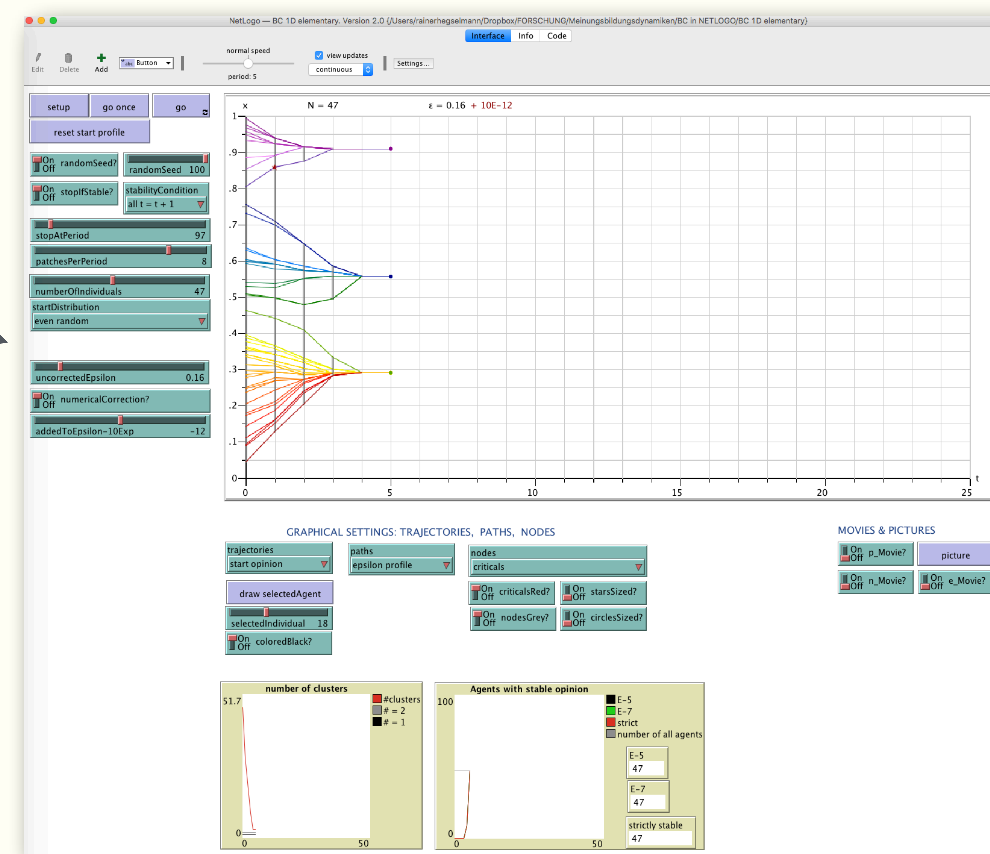

§1
The bounded confidence model: Idea and short analysis

The original starting point: Let's suppose ...

- a group of people, for instance a *group of experts* on something;
- each expert has an *opinion* on the topic under discussion, for instance the probability of a certain type of accident;
- *nobody is totally sure* that he is totally right;
- to some degree everybody is *willing to revise* his opinion when informed about the opinions of others, especially the opinions of *‚competent' others*;
- the revisions produce a new opinion distribution which may lead to further revisions of opinions, and so on and so on.... .

De Vergadering (The meeting), Willy Belinfante

- [Delphi study format / peer disagreement]

The formal definition of the *BC-model*

Each individual takes seriously only those others whose opinions are ‚reasonable‘, ‚not too strange‘, i.e. not too far away from one's own opinion.

FORMAL STATUS

- There is a set of n *individuals; $i, j \in I$.* — set
- Time is *discrete;* $t = 0, 1, 2, \ldots$. — set
- Each individual starts with a certain *opinion*, given by a *real number;* $x_i(0) \in [0,1]$.
- The *profile* of all opinions at time t is — sequence of profiles

$$X(t) = x_1(t), \ldots, x_i(t), x_j(t), \ldots, x_n(t).$$

- For an updating of opinions, each individual i takes into account only those individuals whose opinions are not too far away, i.e. for which $|x_i(t) - x_j(t)| \le \epsilon$ (*confidence level, confidence interval*). — variable

  The *set* of all others that i takes into account at time t is:

„ε–insiders“ → $I(i, X(t)) = \{ j \mid |x_i(t) - x_j(t)| \le \epsilon \}$. — set

„ε–outsiders“ all others

- The next period's opinion of individual i is the *average* opinion of all those which i takes into account: — n-dimensional dynamical system

$$x_i(t+1) = \frac{1}{\#\left(I\left(i, X(t)\right)\right)} \sum_{j \in I\left(i, X(t)\right)} x_j(t)$$

■ : *Red words/phrases give already some vague interpretation of the formalism and hint to the type of intended applications.*

BC: *One simple* formalism with *many* interpretations …

OPINIONS:

- probabilities / degrees of belief for any quantitative or qualitative proposition
- any real-valued quantitative propositions (the normalised range [0,1] does not matter).
- intensity or importance of a wish or preference (*iff* intersubjectively comparable!)
- moral praiseworthiness (0: extremely bad, 0.5: neutral, 1: extremely good)
- budget share

CONTEXTS:

Peer disagreement [Delphi]

Compromise

For reasons as uncertainty, respect for others, an interest in a compromise, a preference for conformity, or due to some social pressure, everybody is willing to compromise with others–but there are limits.

Confirmation bias

The formal definition of the *BC model*

Each individual takes seriously only those others whose opinions are ‚reasonable', not too strange', i.e. not too far away from one's own opinion.

FORMAL STATUS

- There is a set of n *individuals*; $i, j \in I$. — set
- Time is *discrete*; $t = 0, 1, 2, \ldots$. — set
- Each individual starts with a certain *opinion*, given by a *real number*; $x_i(0) \in [0,1]$.
- The *profile* of all opinions at time t is

$$X(t) = x_1(t), \ldots, x_i(t), x_j(t), \ldots, x_n(t).$$ — sequence of profiles

- For an updating of opinions, each individual i takes into account only those individuals whose opinions are not too far away, i.e. for which $|x_i(t) - x_j(t)| \leq \varepsilon$ (*confidence level, confidence interval*). — variable

  The *set* of all others that i takes into account at time t is:

$$I(i,X(t)) = \{ j \mid |x_i(t) - x_j(t)| \leq \varepsilon \}.$$ — set

- The next period's opinion of individual i is the *average* opinion of all those which i takes into account:

$$x_i(t+1) = \frac{1}{\#(I(i,X(t)))} \sum_{j \in I(i,X(t))} x_j(t)$$ — n-dimensional dynamical system

■: Red words/phrases give already some vague interpretation of the formalism.

Centralised social media

A central algorithmic coordination brings together users whose opinions are not too far away from one another.

Decentralised social media

Users send their opinion to all others. Receivers take into account only opinions that are not too far away.

PURPOSES / STATUS:

- descriptive
- normative
- technical

Some examples:
$n = 50$, *same* even random start distribution

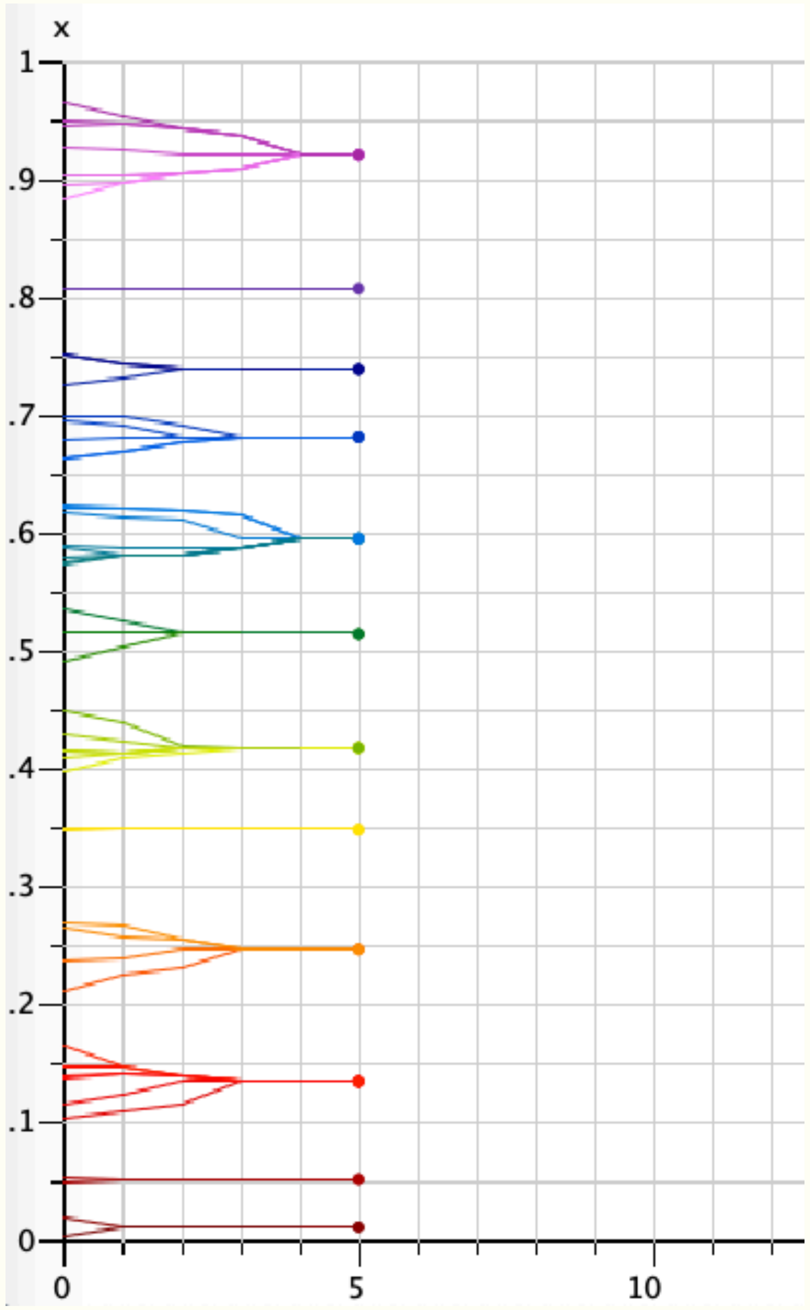

$\epsilon = 0.03$

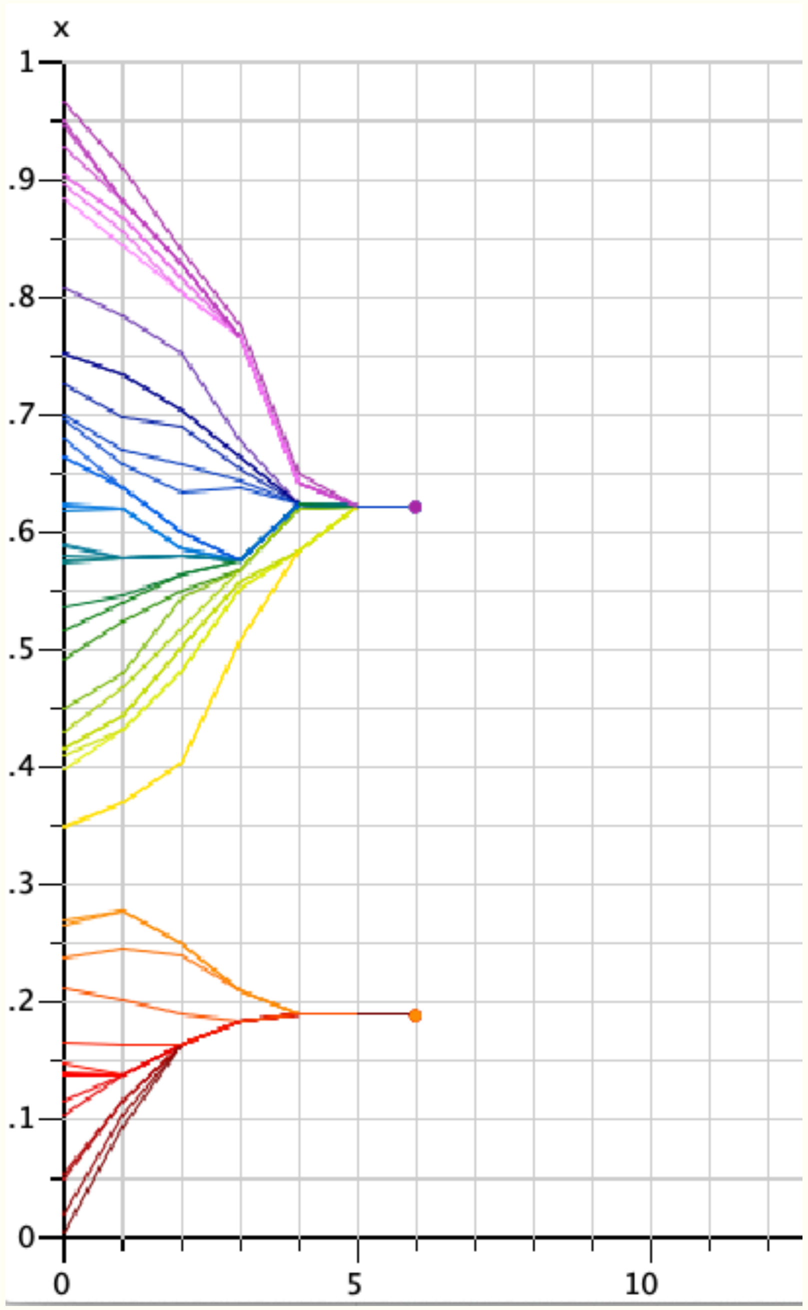

$\epsilon = 0.2$

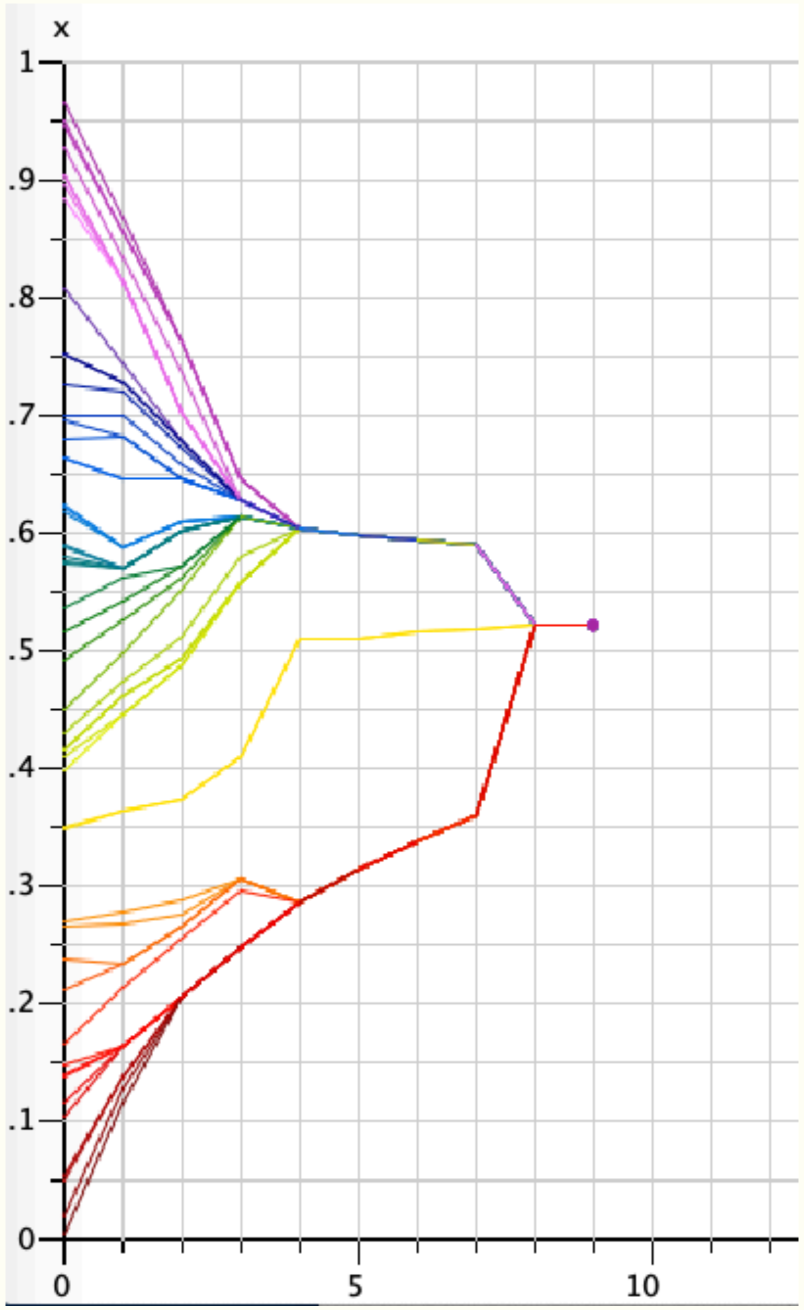

$\epsilon = 0.25$

Understanding the BC-model
Some terminology: ϵ–profiles

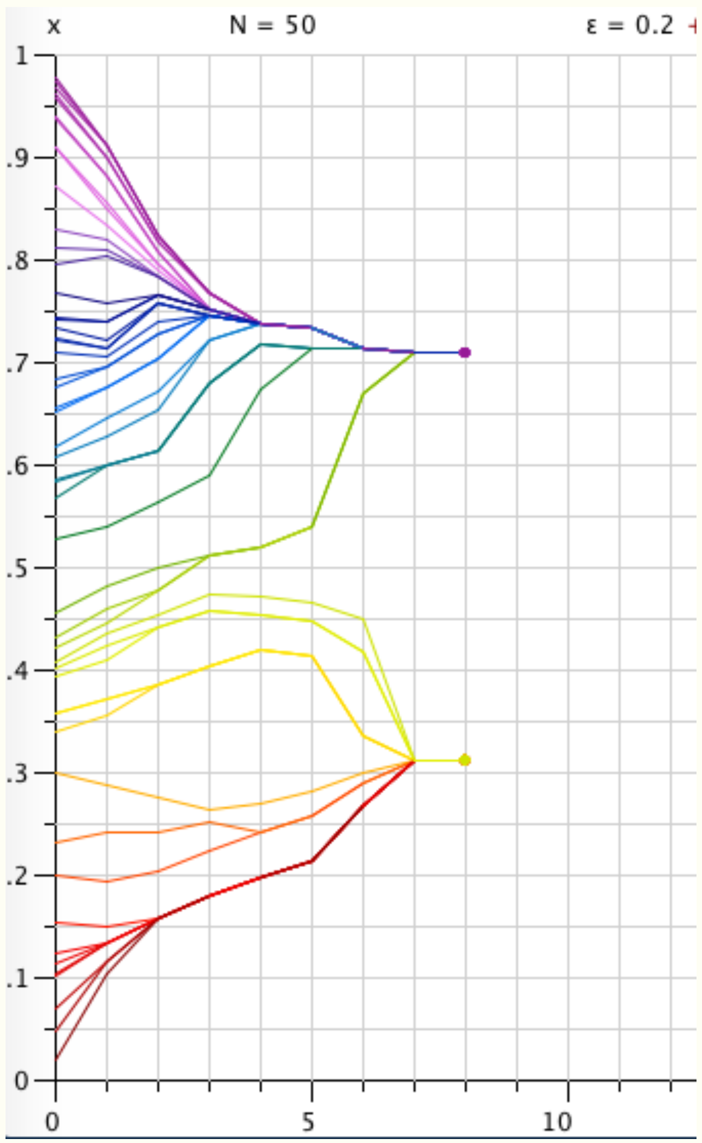

<u>Definition 1</u>

The opinion profile

$x(t) = x_1(t), x_2(t), \ldots, x_i(t), \ldots x_n(t)$

is an *ordered opinion profile* iff

$0 \leq x_1(t) \leq x_2(t) \leq \ldots \leq x_i(t) \leq \ldots \leq x_n(t)$

<u>Definition 2:</u>

An ordered opinion profile is an *ε–profile* iff for all $i = 2, \ldots, n$ it holds

$x_{i+1} - x_i \leq \epsilon.$

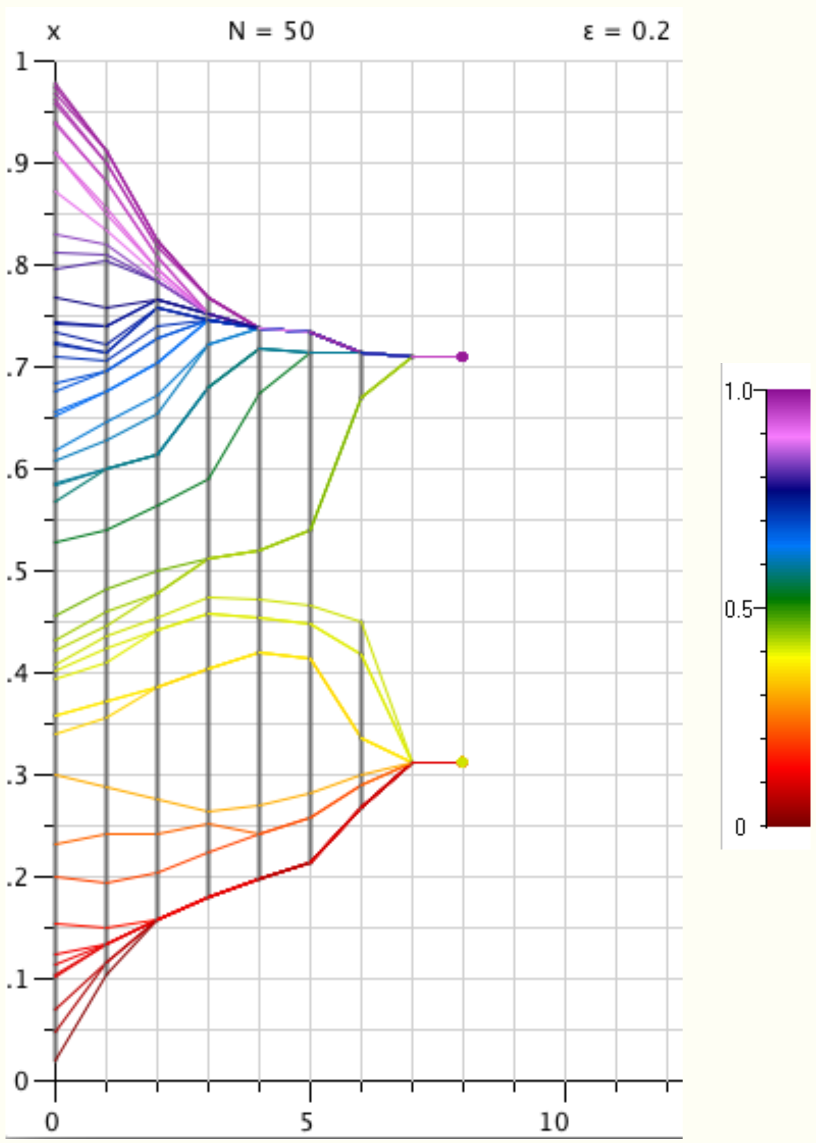

Understanding the BC-model
A key concept: the ϵ–split

Extreme opinions are under a *one sided influence* and move direction centre. The range of the profile *shrinks*.

At the extremes opinions *condense*.

Condensed regions *attract* opinions from less populated areas within their ϵ–reach. In the centre opinions > 0.5 move *upwards*, opinions < 0.5 *downwards*.

The ε-profile *splits* in t_4. From now on the split sub-profiles belong to different ‚opinion worlds‘ or communities which *do no longer interact*.

In the two split off sub-profiles, in t_5 *all* opinions have *all* opinions within their confidence interval. Therefore, in t_6 all opinions *collapse* into one and the same opinion.

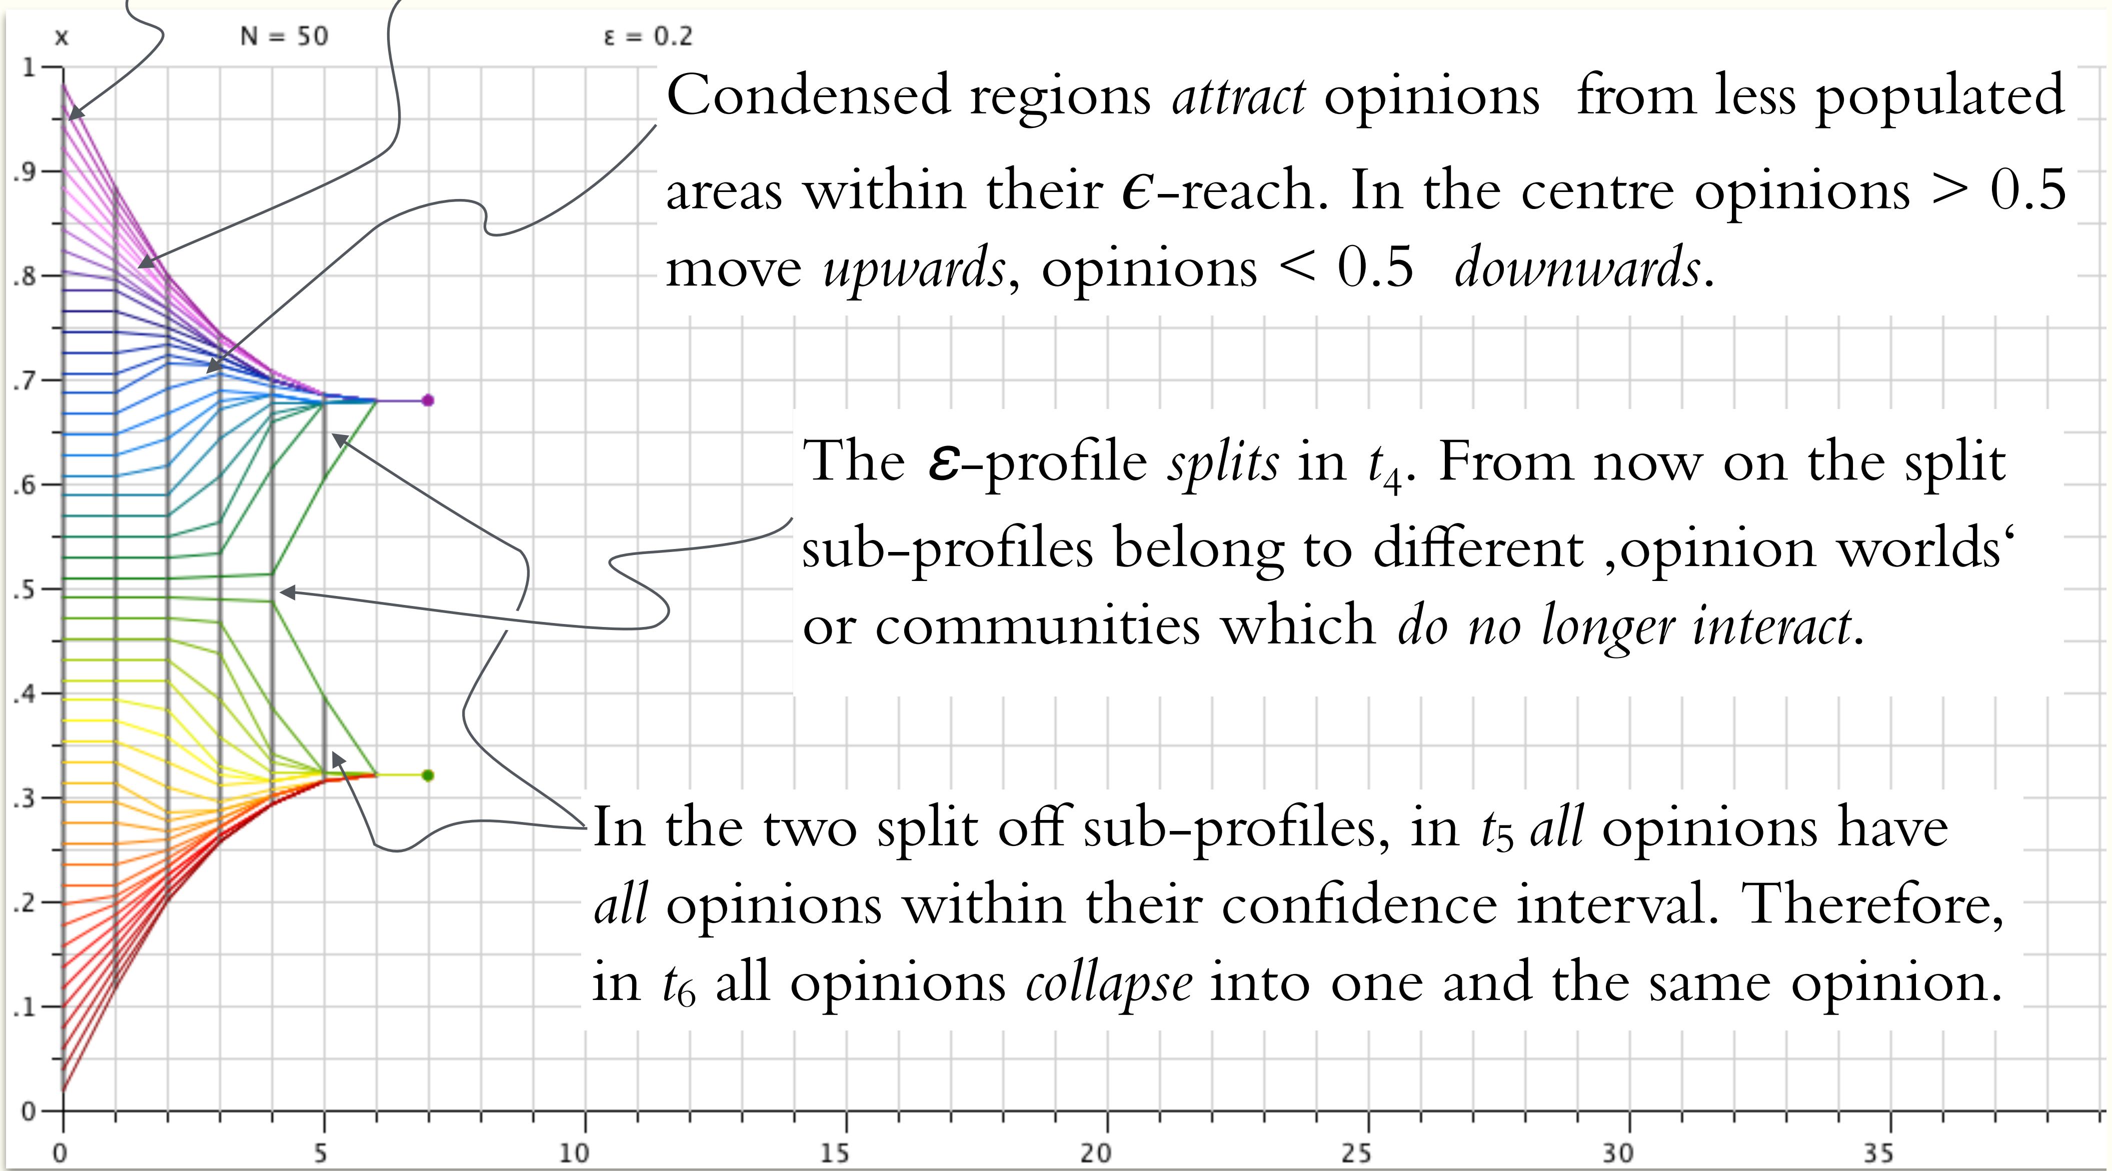

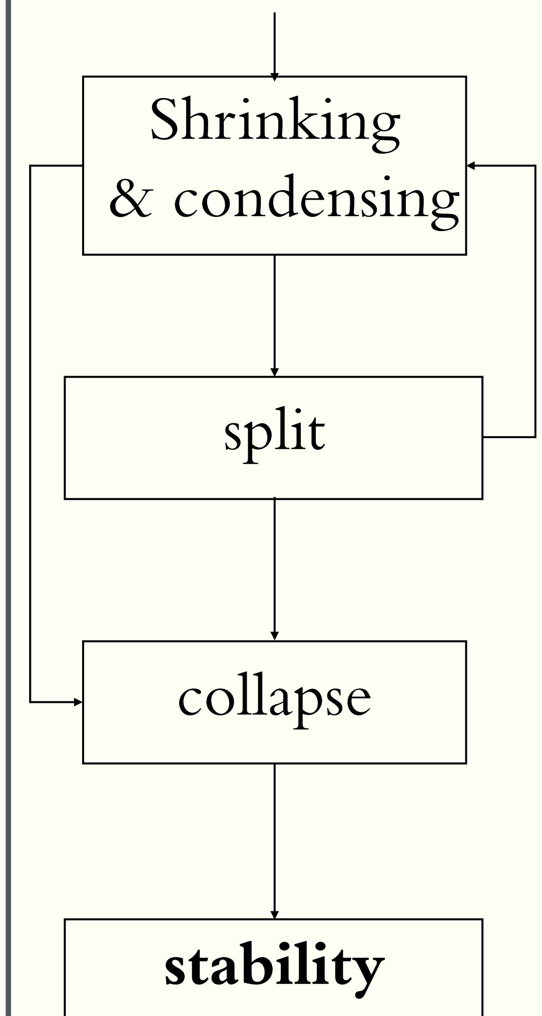

Dynamics with 50 opinions, simultaneous updating, equidistant start profile, ε = 0.2.

Old view$_{\text{HK 2002}}$: Polarisation as a certain phase

- random start distribution
- 625 opinions
- simultaneous updating
- 41 steps of increasing confidence: $\varepsilon = 0, 0.01, \ldots, 0.4$;
- 50 repetitions per step, each running until stabilised;

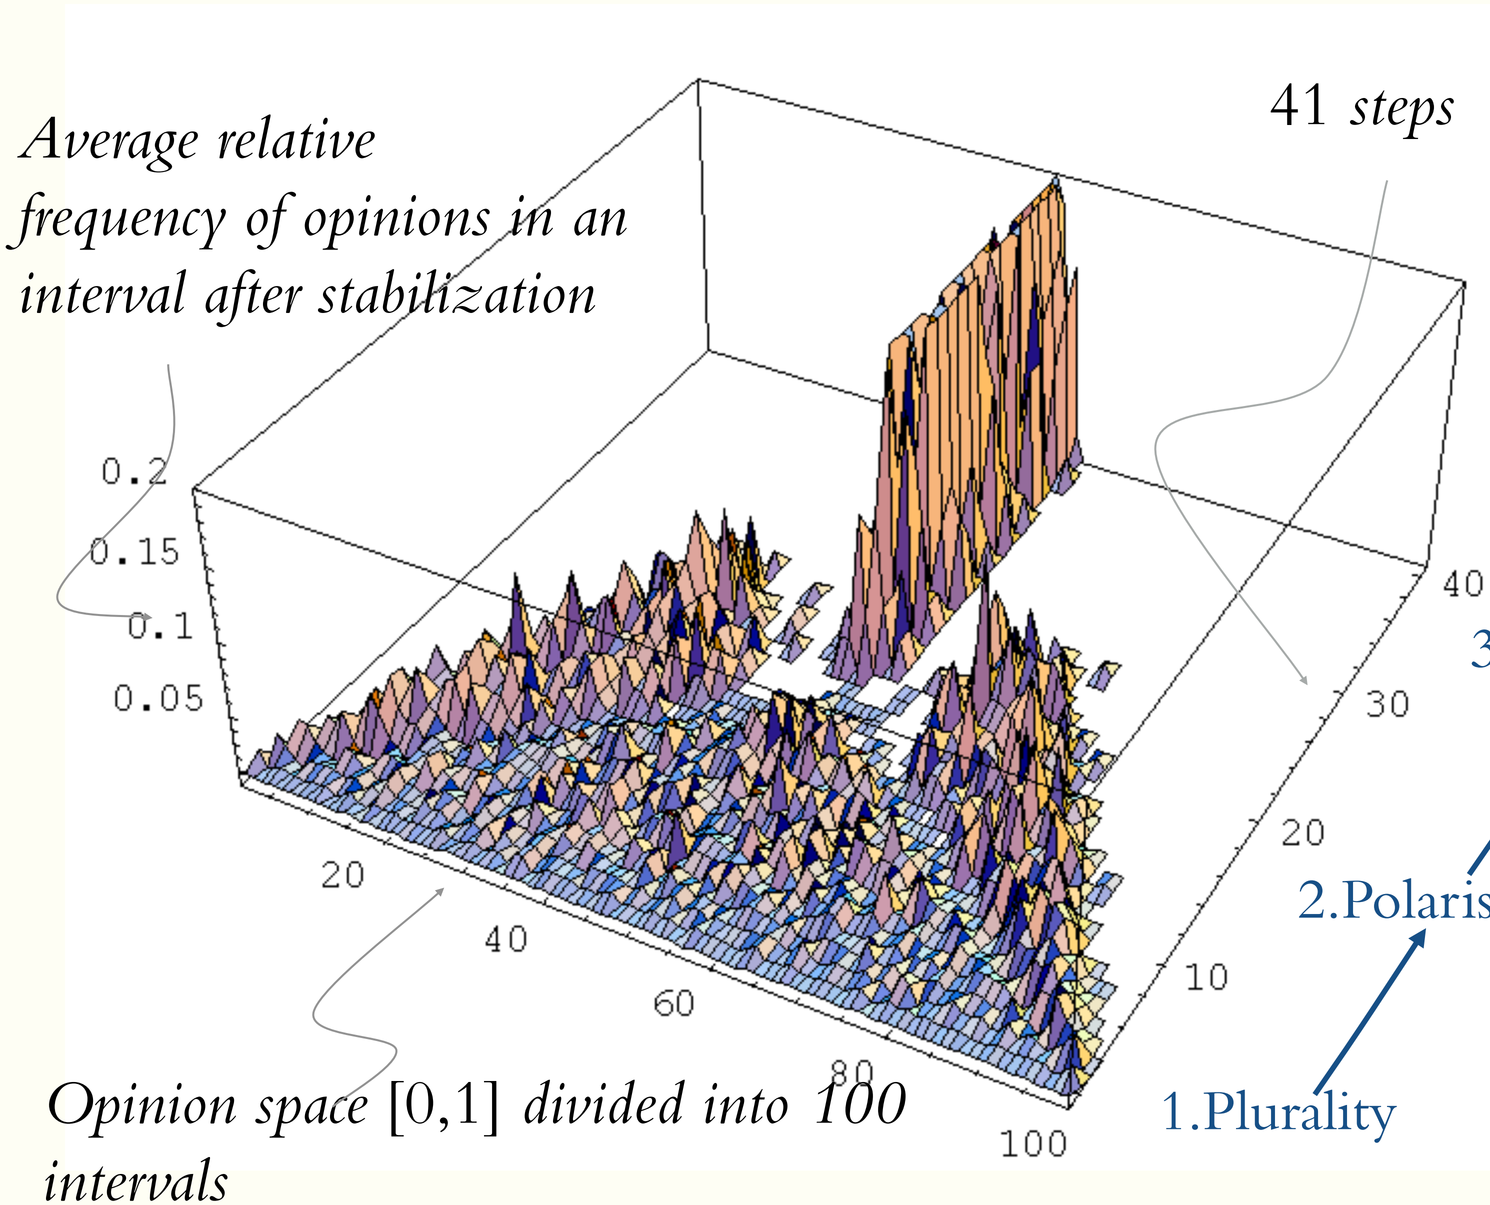

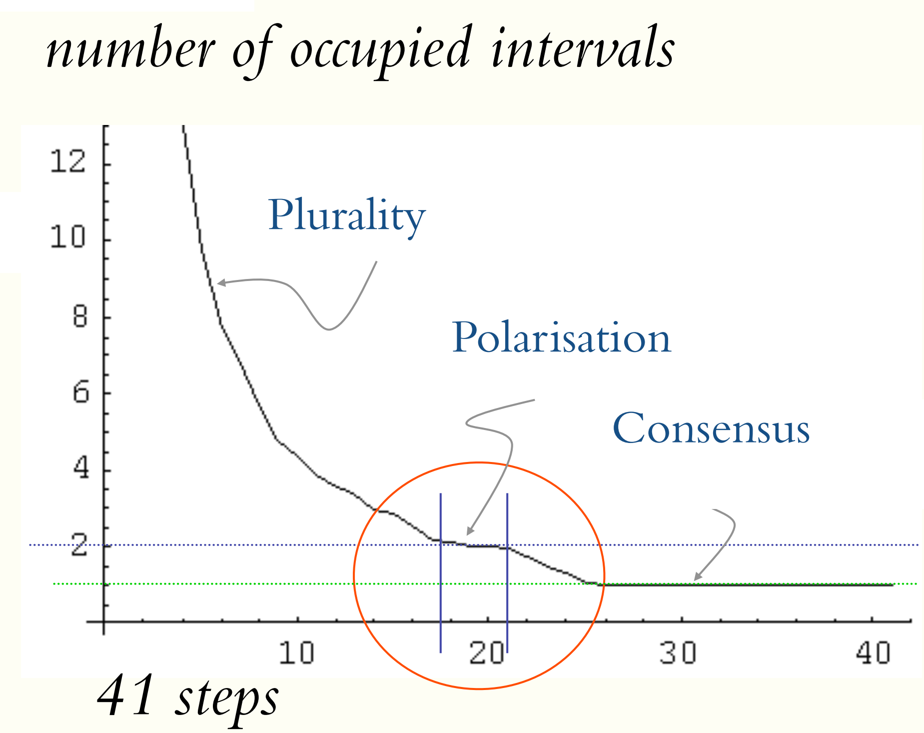

Method: Random start distributions, iterations, averaging

As the confidence level ϵ increases ….

Trivially, it holds:

- For a confidence level $\epsilon = 0$, nothing happens: $X(0) = X(1)$, i.e. stability for $t = 1$.
- For a confidence level $\epsilon = 1$, for *all* agents i their opinion $x_i(1)$ is the mean of all opinions in $X(0)$, i.e. *consensus*. Furthermore, we get $X(1) = X(2)$, is e. stability for $t = 2$.

Some examples:
$N = 50$, *same* even random start distribution

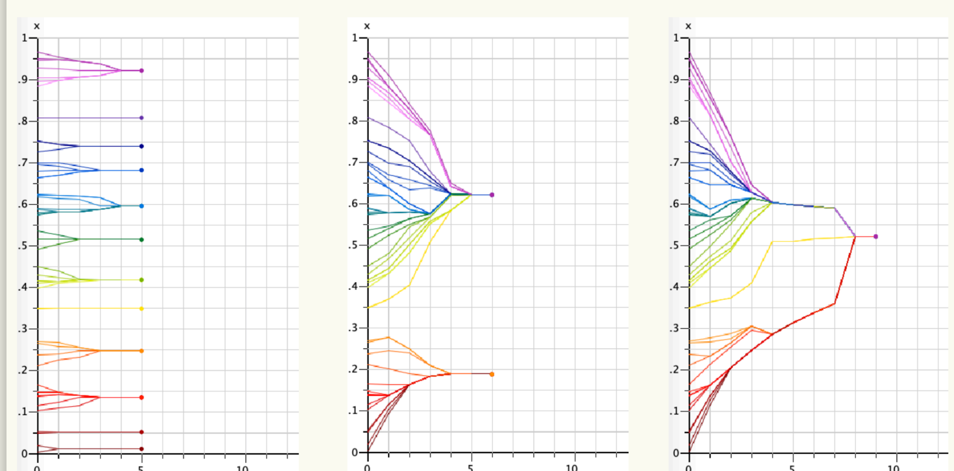

$\epsilon = 0.03$ $\epsilon = 0.2$ $\epsilon = 0.25$

But what's about $0 < \epsilon < 1$?

Suggested by your analysis, but wrong:

- With an *increasing* confidence level ϵ, the number of final clusters *decreases monotonically* from plurality to polarisation to consensus.
- With an *increasing* confidence level ϵ, the width of the final stabilised profile *shrinks monotonically*.

§ 2

A new analytical approach: Given a certain start distribution X(0), how many *different* BC-processes $\langle X(0), \epsilon \rangle$ do exist?

The subject to understand: a *dynamical* network

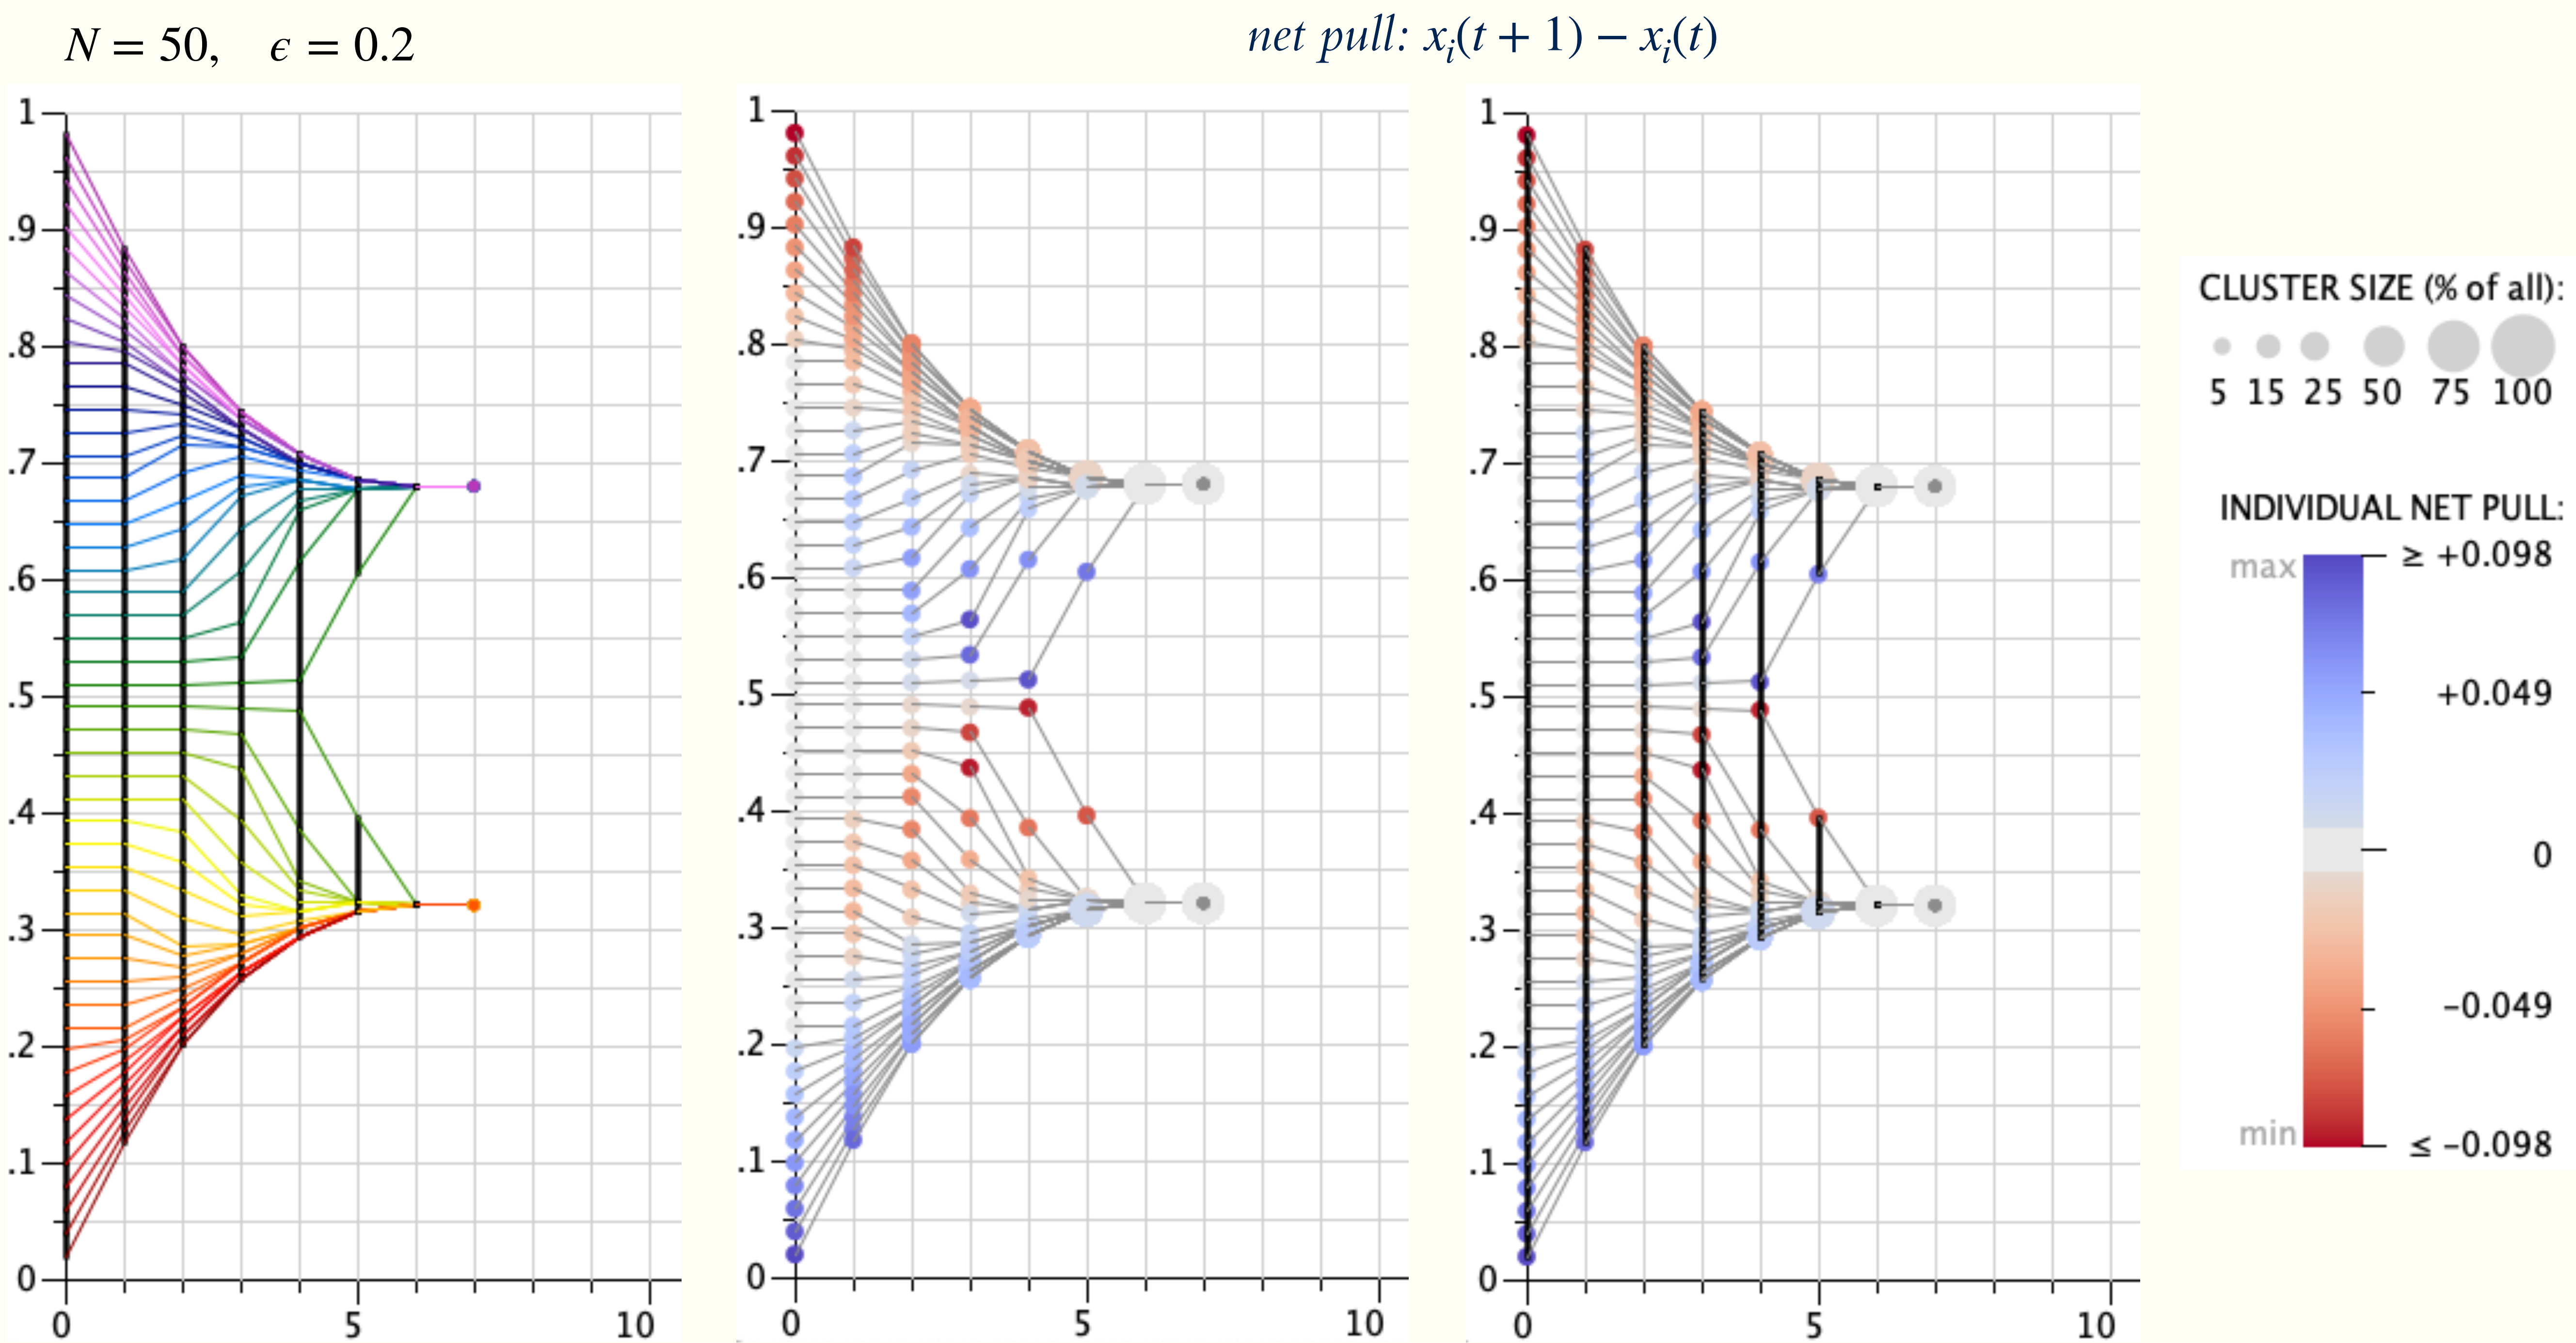

- Vertical lines connect nodes in a network.
- They are paths through nodes that are directly or indirectly linked.
- ϵ-splits indicate that upward and downward forces have torn apart a former network.

Description of the BC model *versus* Understanding effects in the BC model

BC model

formal description in terms of *dynamical systems*

explanation of effects in terms of *network theory*

The formal definition of the *BC model*

Each individual takes seriously only those others whose opinions are ‚reasonable', not too strange', i.e. not too far away from one's own opinion.

FORMAL STATUS

- There is a set of n *individuals*; $i, j \in I$. — set
- Time is *discrete;* $t = 0, 1, 2, \ldots$. — set
- Each individual starts with a certain *opinion*, given by a *real number*; $x_i(0) \in [0,1]$.
- The *profile* of all opinions at time t is — sequence of profiles

$$X(t) = x_1(t), \ldots, x_i(t), x_j(t), \ldots, x_n(t).$$

- For an updating of opinions, each individual i takes into account only those individuals whose opinions are not too far away, i.e. for which $|x_i(t) - x_j(t)| \leq \varepsilon$ (*confidence level, confidence interval*). — variable

  The *set* of all others that i takes into account at time t is: — set

$$I(i,X(t)) = \{j \mid |x_i(t) - x_j(t)| \leq \varepsilon\}.$$

- The next period's opinion of individual i is the *average* opinion of all those which i takes into account: — n-dimensional dynamical system

$$x_i(t+1) = \frac{1}{\#(I(i,X(t)))} \sum_{j \in I(i,X(t))} x_j(t)$$

■ : Red words/phrases give already some vague interpretation of the formalism.

The subject to understand: a *dynamical* network

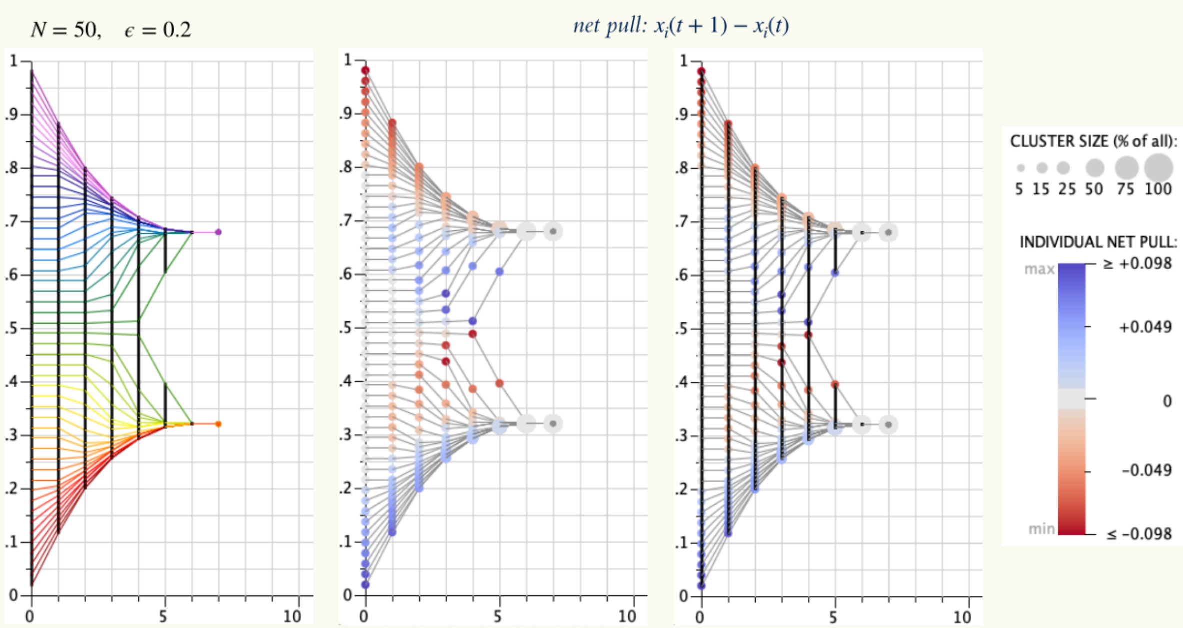

- Vertical lines connect nodes in a network.
- They are paths through nodes that are directly or indirectly linked.
- ϵ-splits indicate that upward and downward forces have torn apart a former network.

As the confidence level ϵ increases ….

Example

Start distribution with 5 agents/opinions: $X(0) = \langle 0.0, 0.18, 0.36, 0.68, 1.0 \rangle$

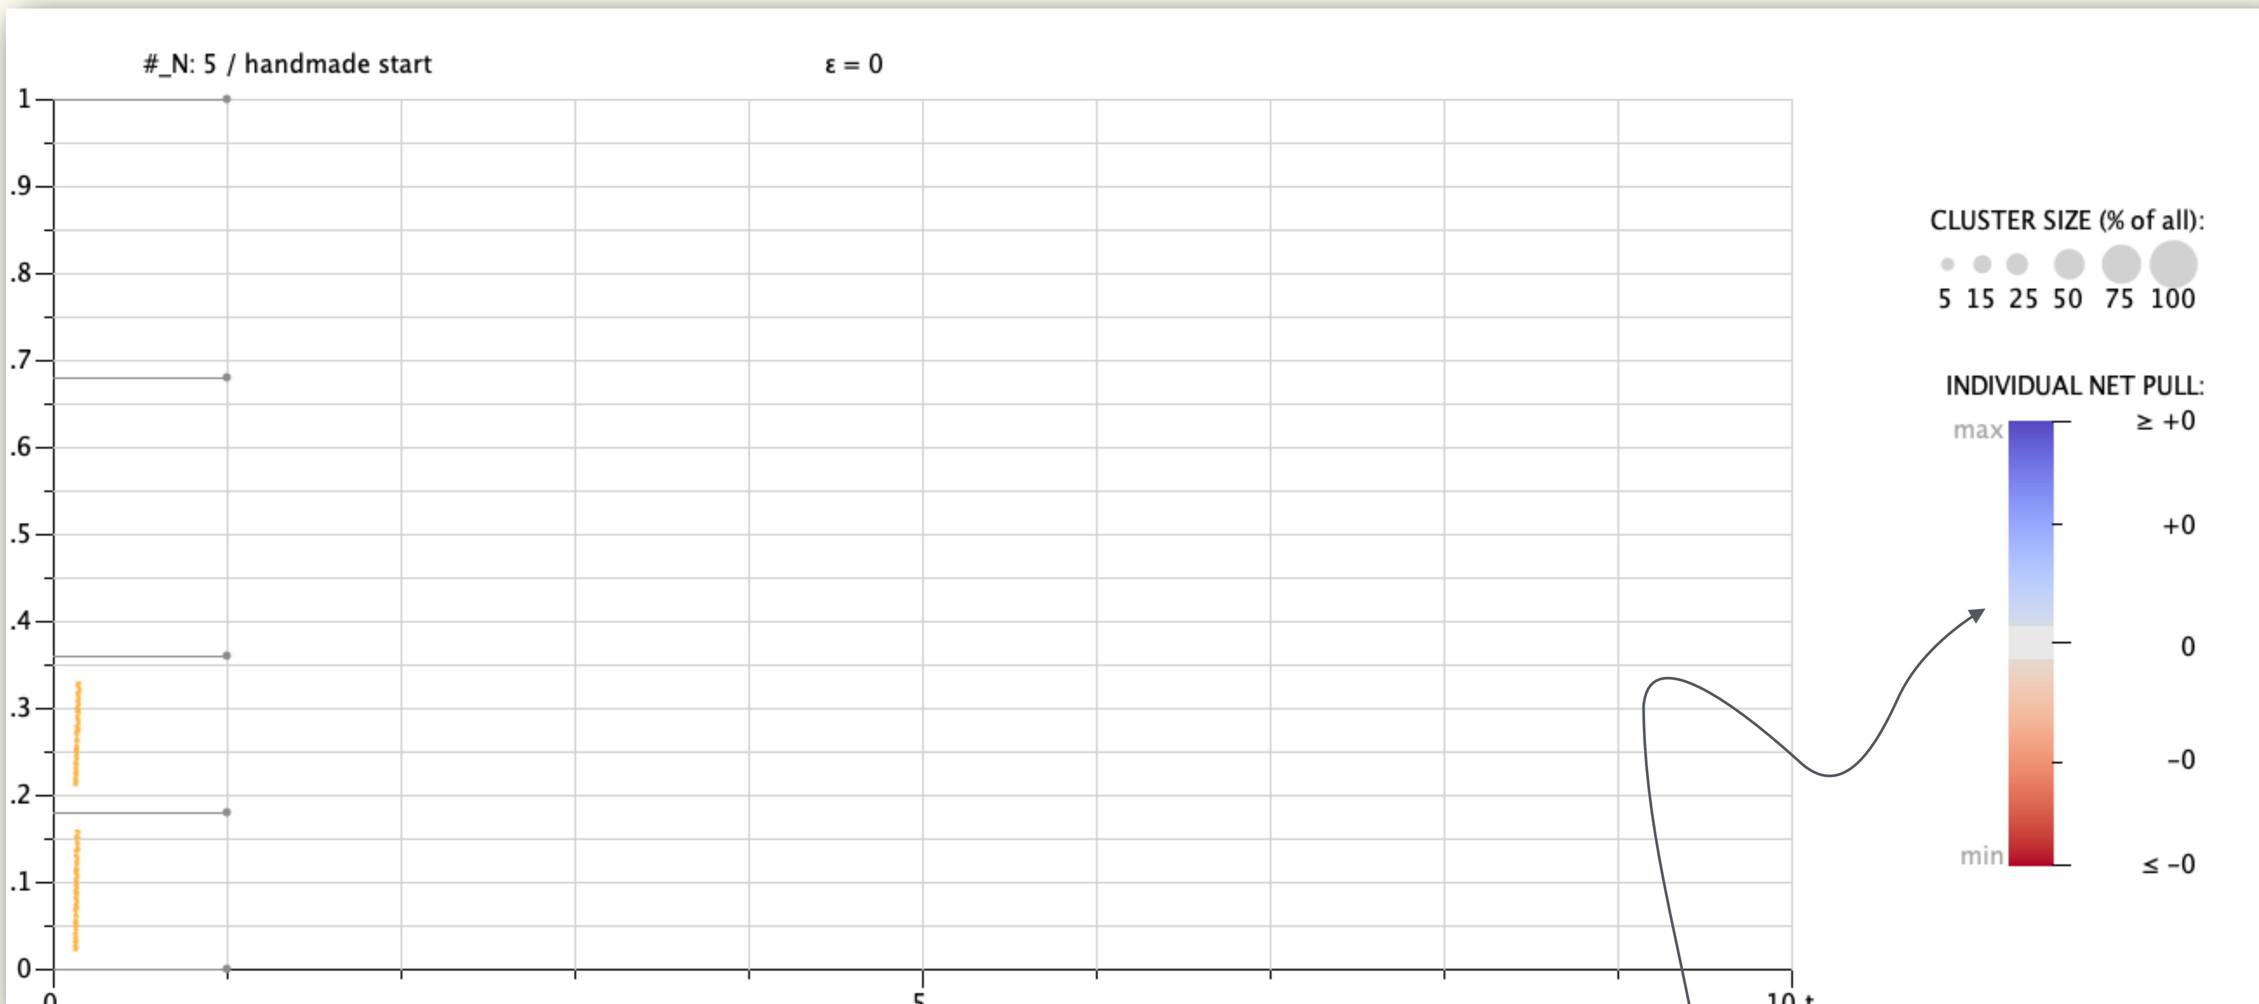

Individual net pull = ($\text{pull}_{\text{upward}}$ – $\text{pull}_{\text{downward}}$) on $x_i(t)$

[$\leftrightarrow x_i(t+1) - x_i(t)$]

A new analytical concept: ϵ–switches

Example

Start distribution with 5 agents/opinions: $X(0) = \langle 0.0, 0.18, 0.36, 0.68, 1.0 \rangle$

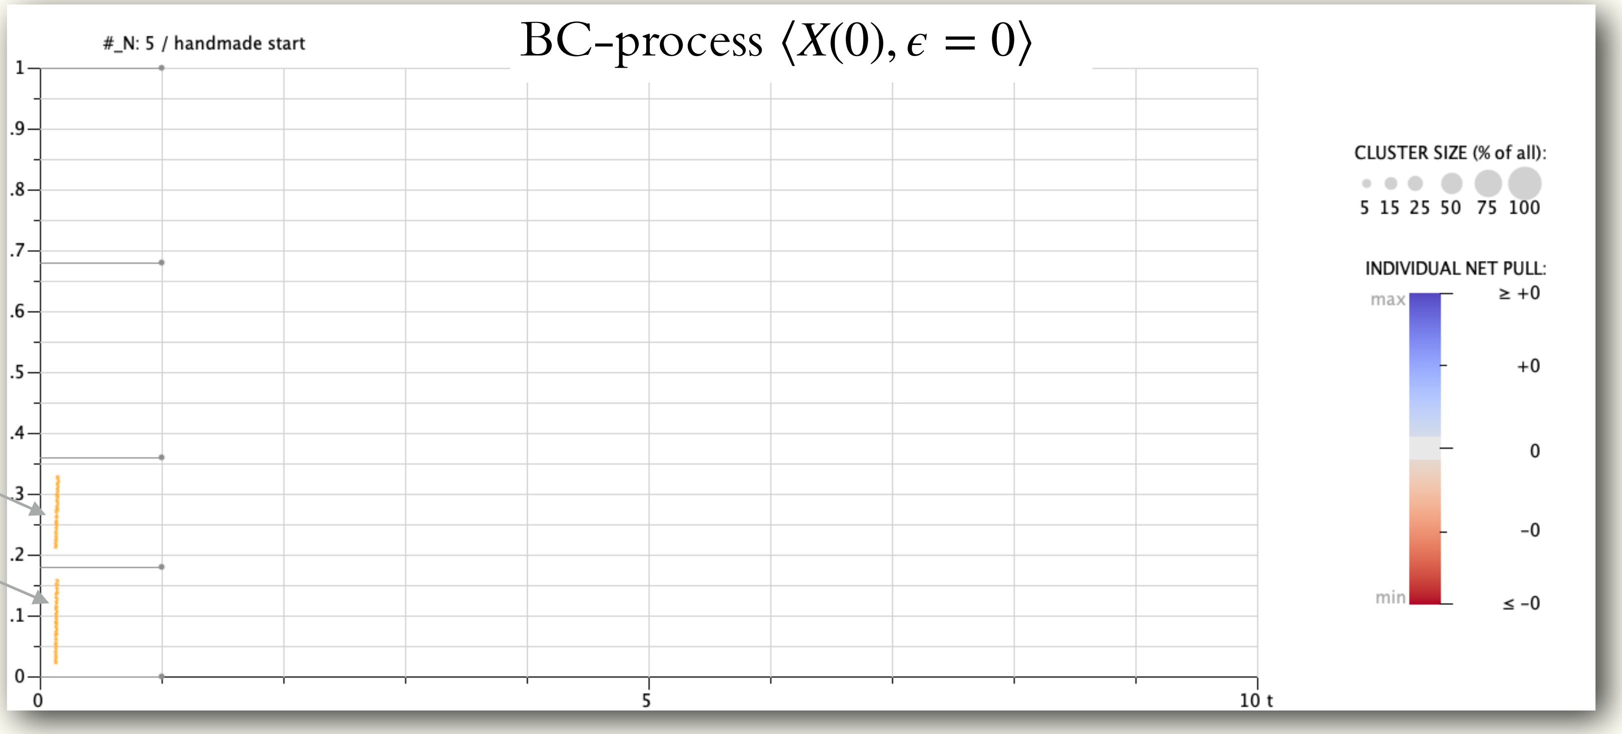

$\epsilon_1^* = 0.18$ is for the dynamics $\langle X(0), \epsilon = 0 \rangle$ an *ϵ-switch:* the smallest *larger* ϵ–value that changes the dynamics $\langle X(0), \epsilon = 0 \rangle$.

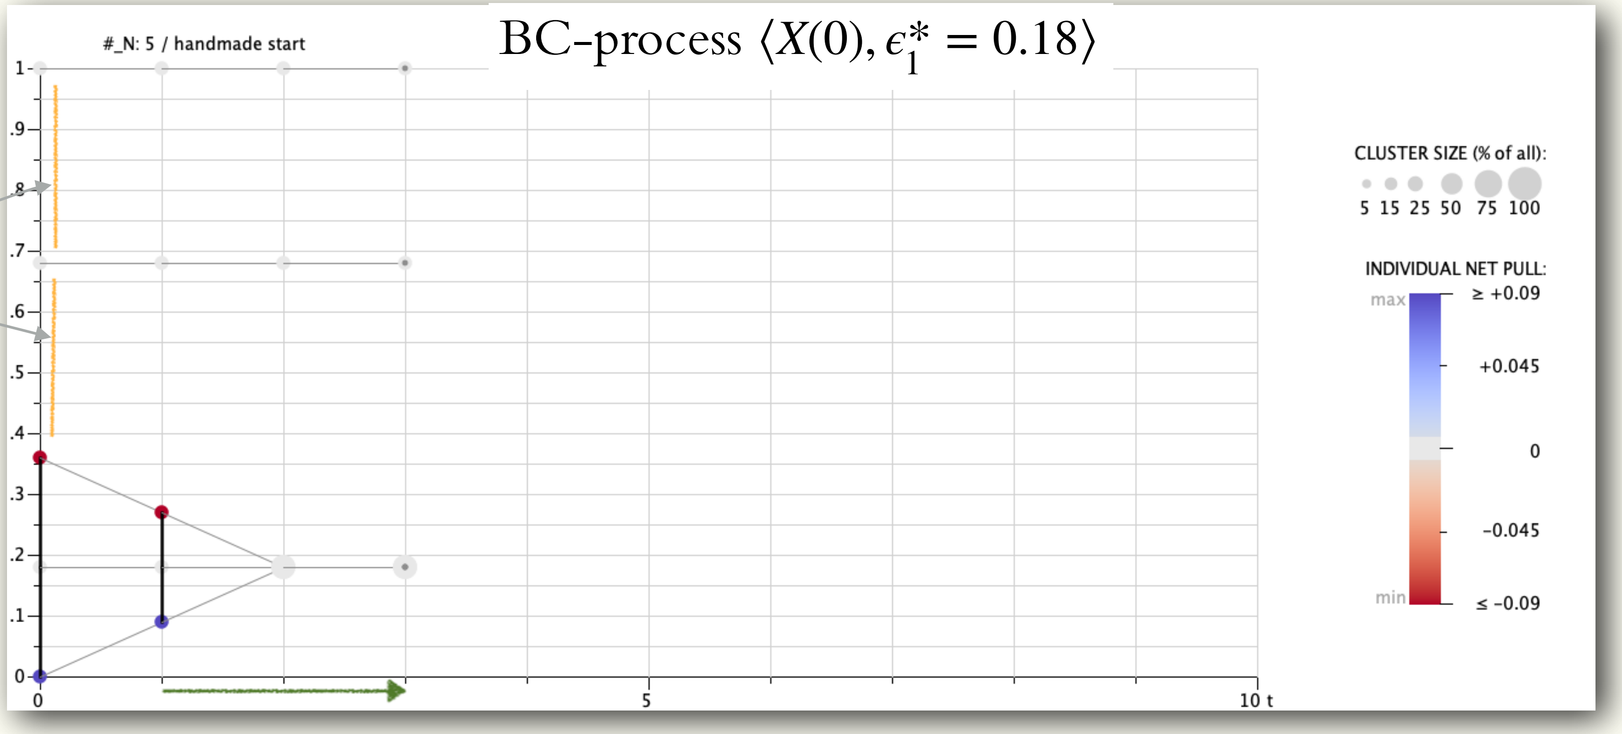

$\epsilon_2^* = 0.32$ is for the dynamics $\langle X(0), \epsilon_1^* = 0.18 \rangle$ an *ϵ-switch:* the smallest *larger* ϵ–value that changes the dynamics $\langle X(0), \epsilon_1^* = 0.18 \rangle$.

A new analytical concept: ϵ–switches

Example

Start distribution with 5 agents/opinions: $X(0) = \langle 0.0, 0.18, 0.36, 0.68, 1.0 \rangle$

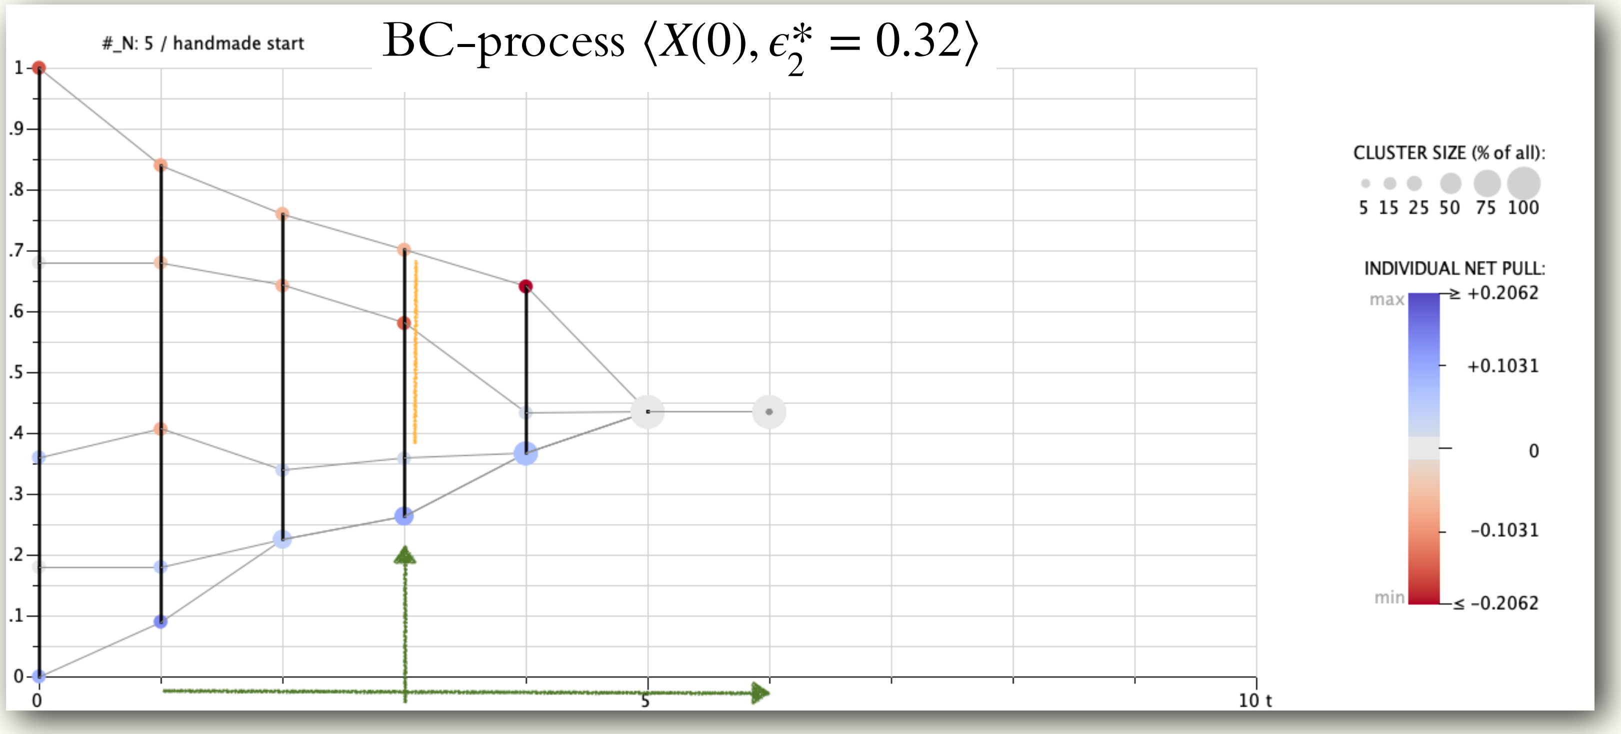

Questions:

What is the next ϵ–switch, i.e. the smallest *larger* ϵ–value that changes the dynamics $\langle X(0), \epsilon_2^* = 0.32 \rangle$?

Answer:

It is the *distance to a nearest ϵ-outsider* that we can find in the whole BC–process $\langle X(0), \epsilon_2^* = 0.32 \rangle$.

$$\delta_{min}^{out}\big(X(0), \epsilon\big) = min\Big\{ |x_i(t) - x_j(t)| \,\Big|\, t = 0, 1, \ldots, \bar{t};\ i = 1, \ldots, n;\ j \in O\big(i, X(t), \epsilon\big) \Big\}$$

Idea: Algorithmic search for all ϵ-switches of a given start distribution $X(0)$

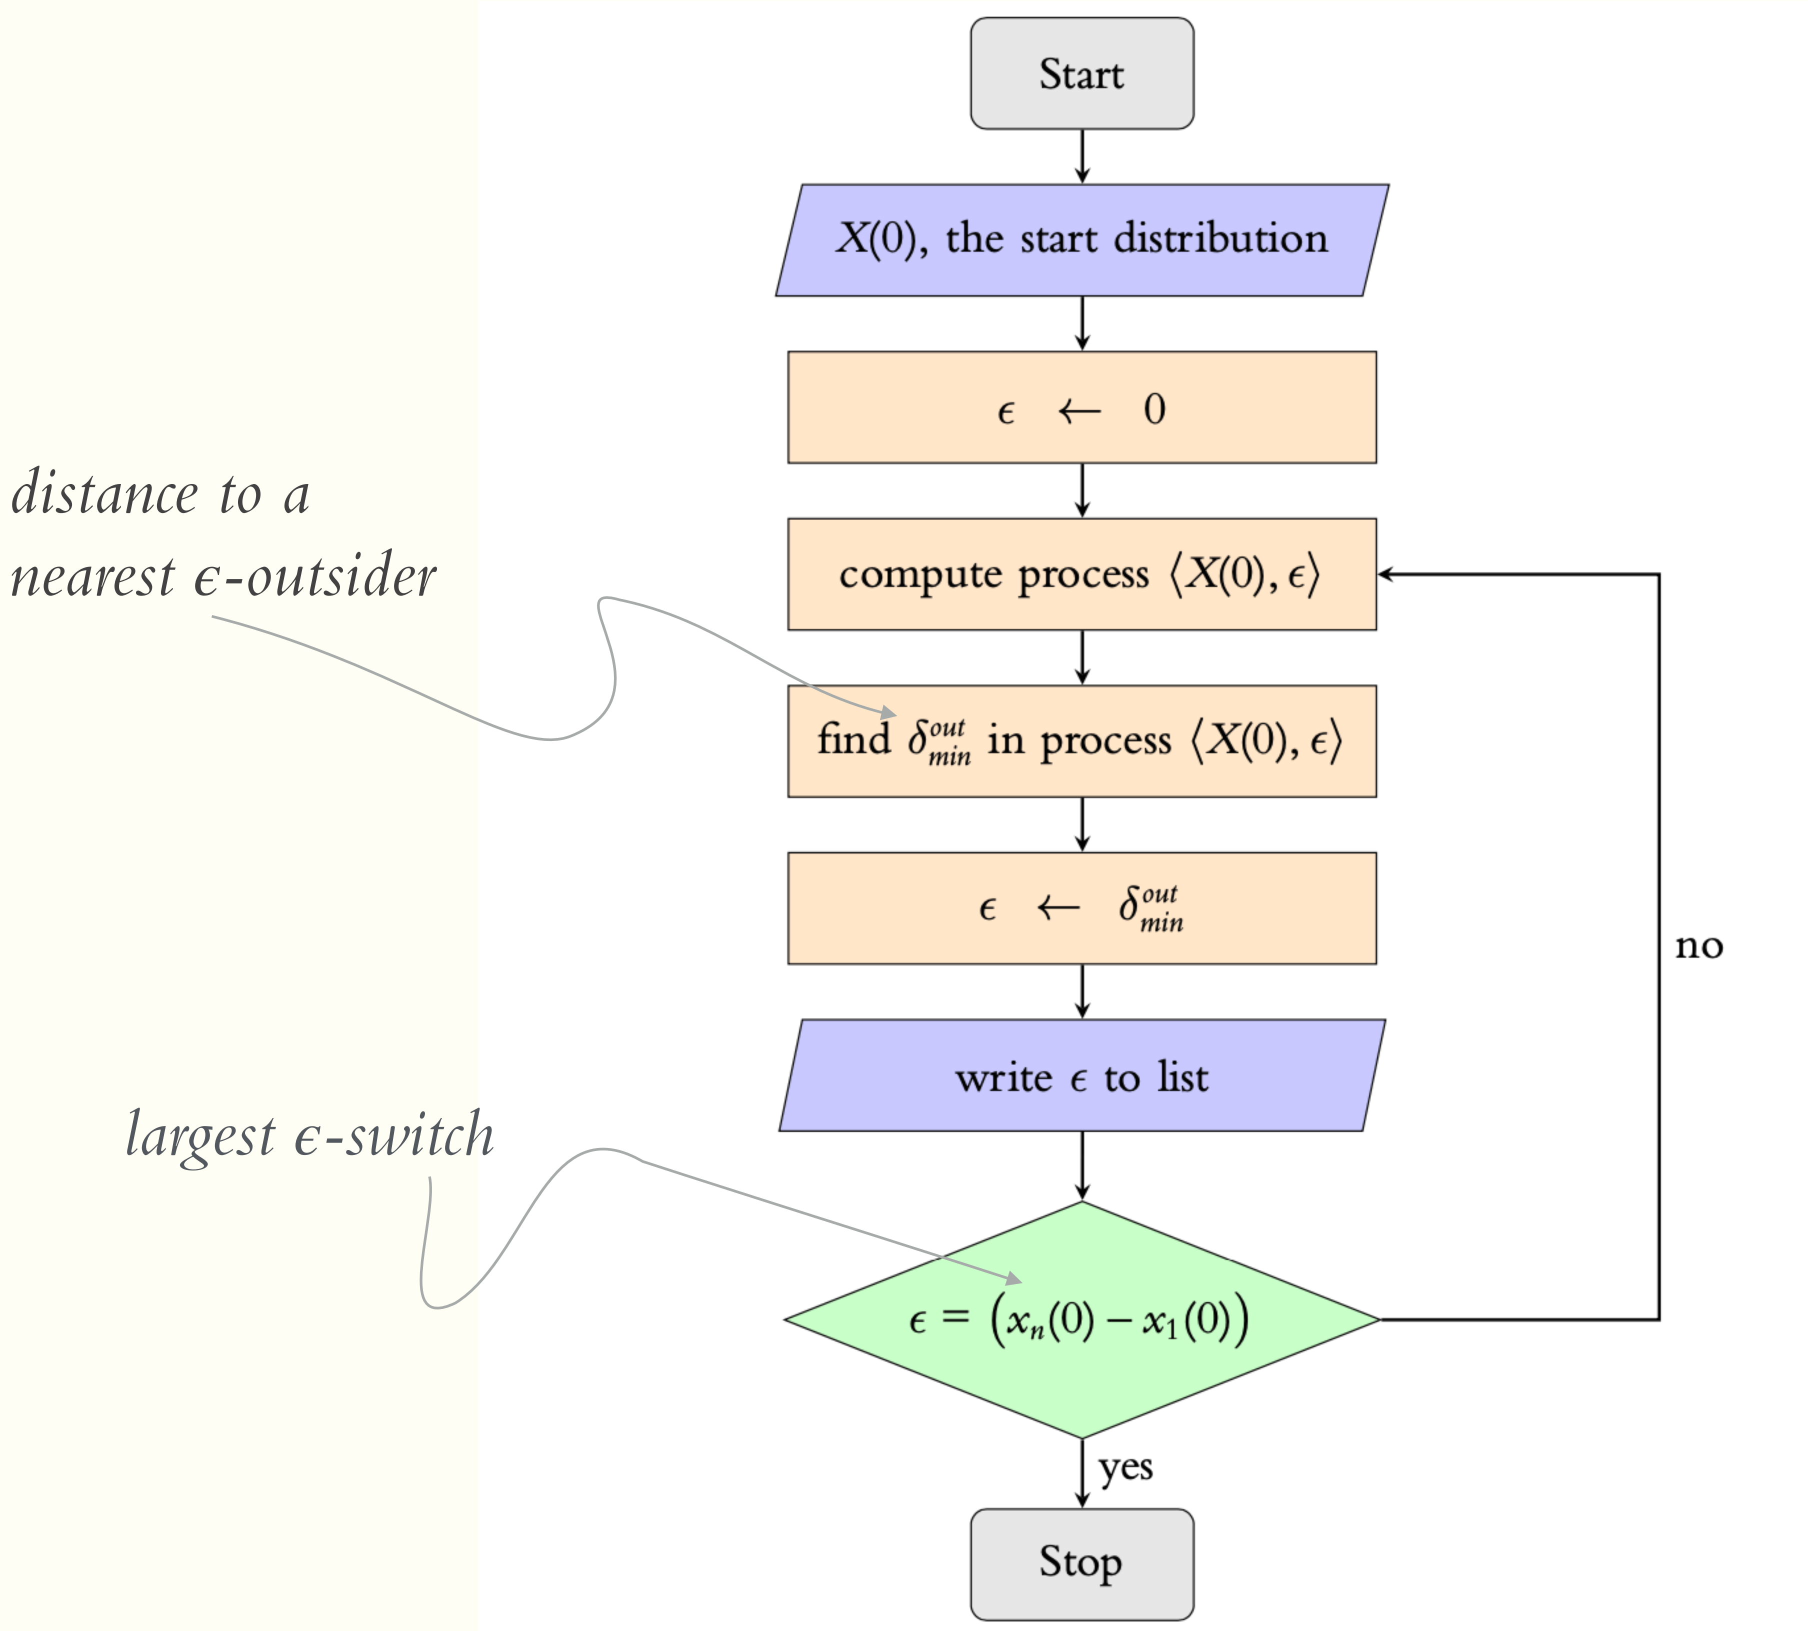

From switch ϵ_3^* to switch ϵ_4^*

Example

Start distribution with 5 agents/opinions: $X(0) = \langle 0.0, 0.18, 0.36, 0.68, 1.0 \rangle$

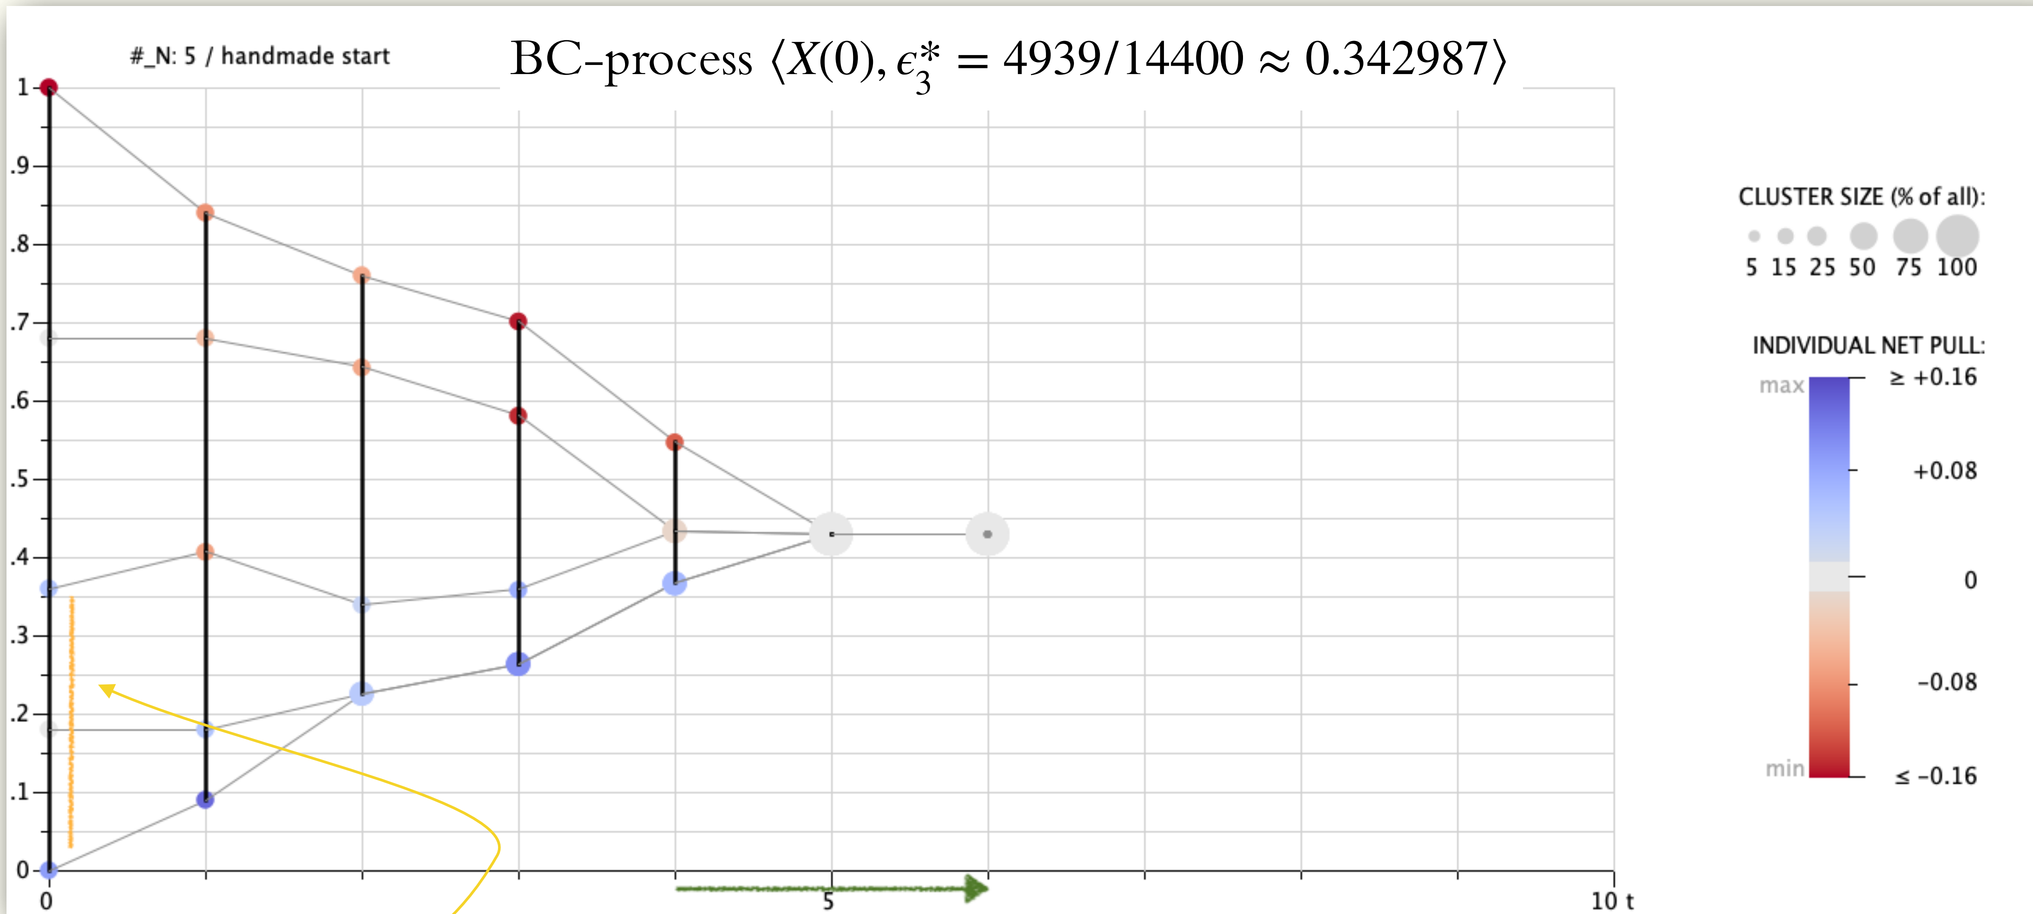

distance to a
nearest ϵ-outsider

From switch ϵ_4^* to switch ϵ_5^*

Example

Start distribution with 5 agents/opinions: $X(0) = \langle 0.0,\ 0.18,\ 0.36,\ 0.68,\ 1.0 \rangle$

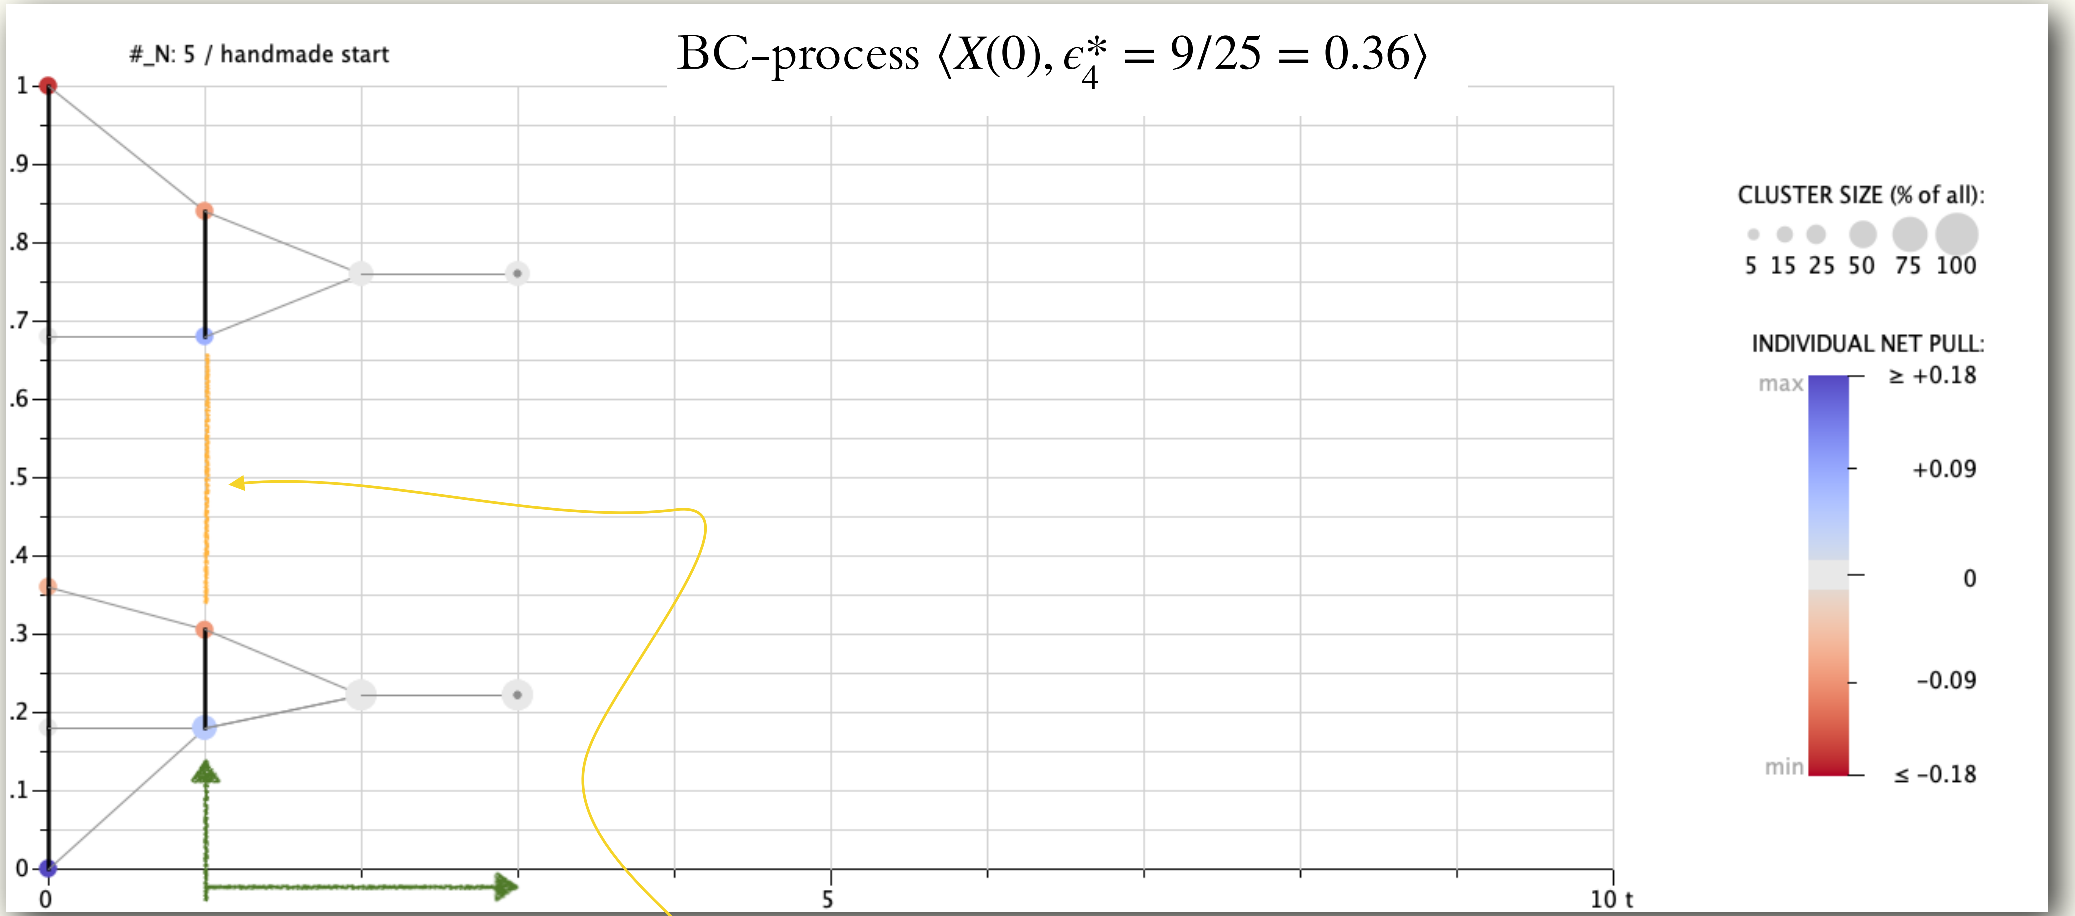

distance to a nearest ϵ-outsider

From switch ϵ_5^* to switch ϵ_6^*

Example

Start distribution with 5 agents/opinions: $X(0) = \langle 0.0, 0.18, 0.36, 0.68, 1.0 \rangle$

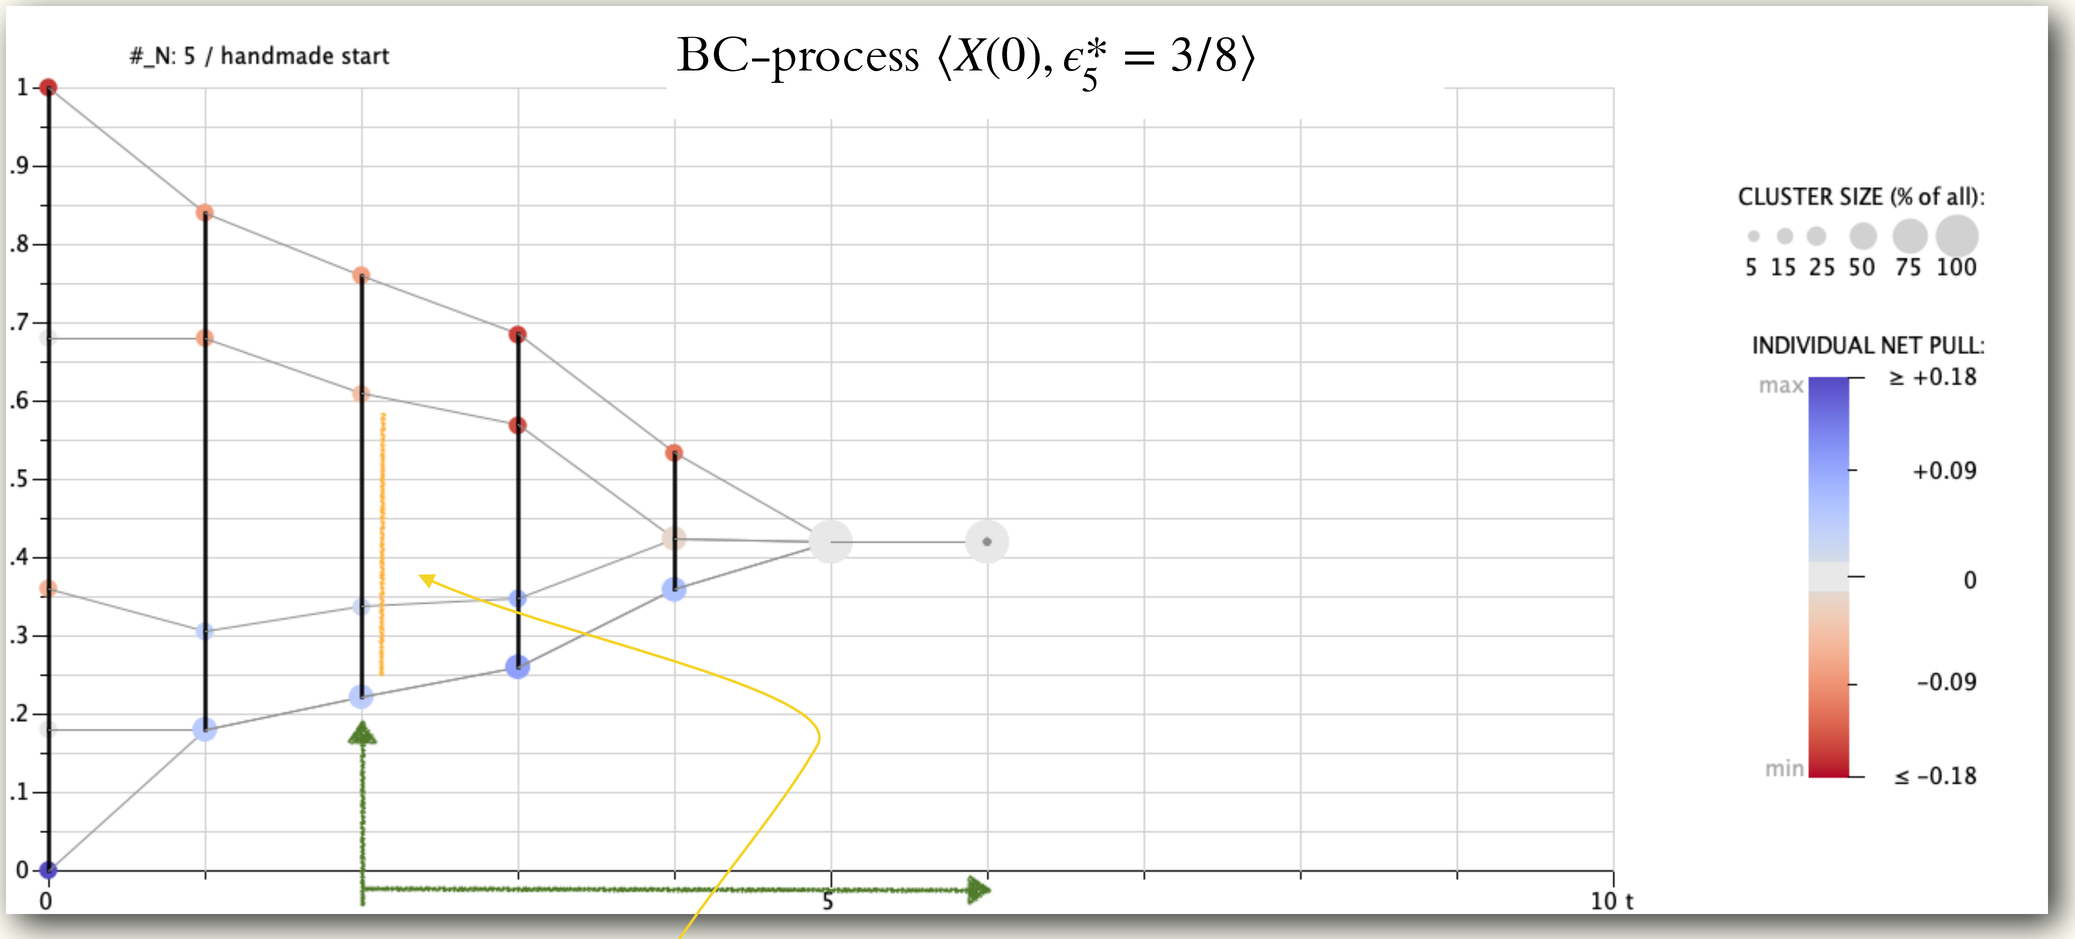

distance to a
nearest ϵ-outsider

In our example:
There are 13 ϵ-switches that make a difference

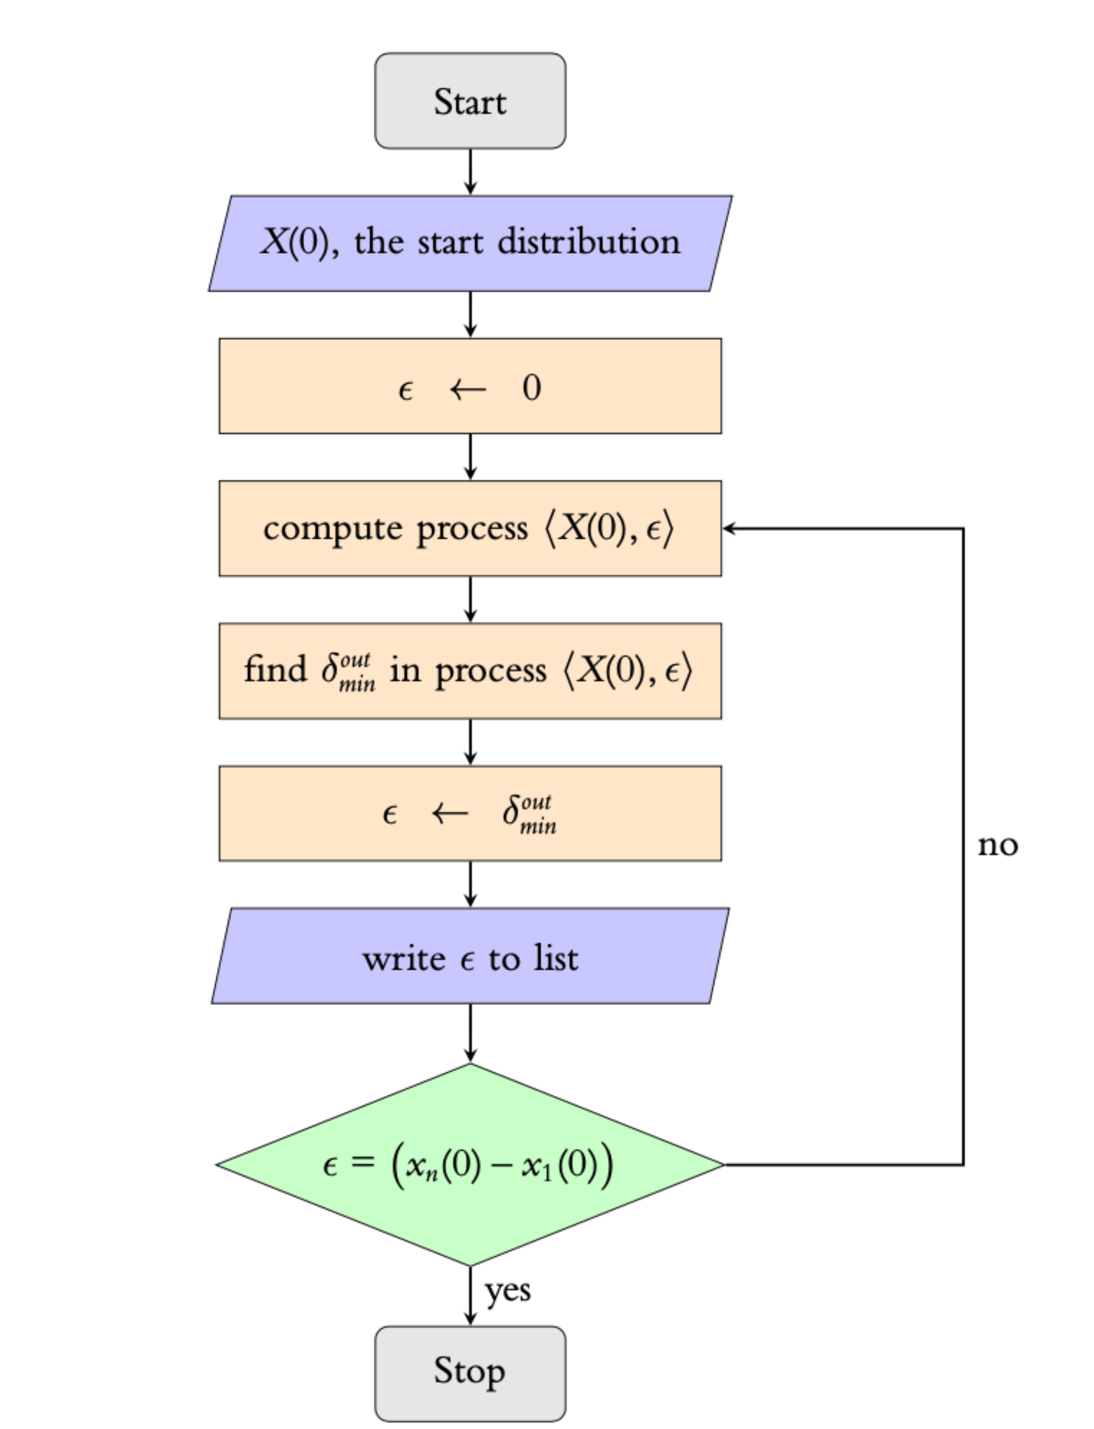

switch	exact	float64
ϵ^*_1	9/50	0.18
ϵ^*_2	8/25	0.32
ϵ^*_3	4939/14400	0, 342986111111111
ϵ^*_4	9/25	0.36
ϵ^*_5	3/8	0.375
ϵ^*_6	29/75	0.386666666666667
ϵ^*_7	339/800	0.42375
ϵ^*_8	1/2	0.5
ϵ^*_9	107/200	0.535
ϵ^*_{10}	16/25	0.64
ϵ^*_{11}	17/25	0.68
ϵ^*_{12}	41/50	0.82
ϵ^*_{13}	1/1	1.0

Excursion: Floating-point arithmetic

BC-processes: For the computer too difficult!

Equidistant start profile with 4 agents:

What the computer calculates:

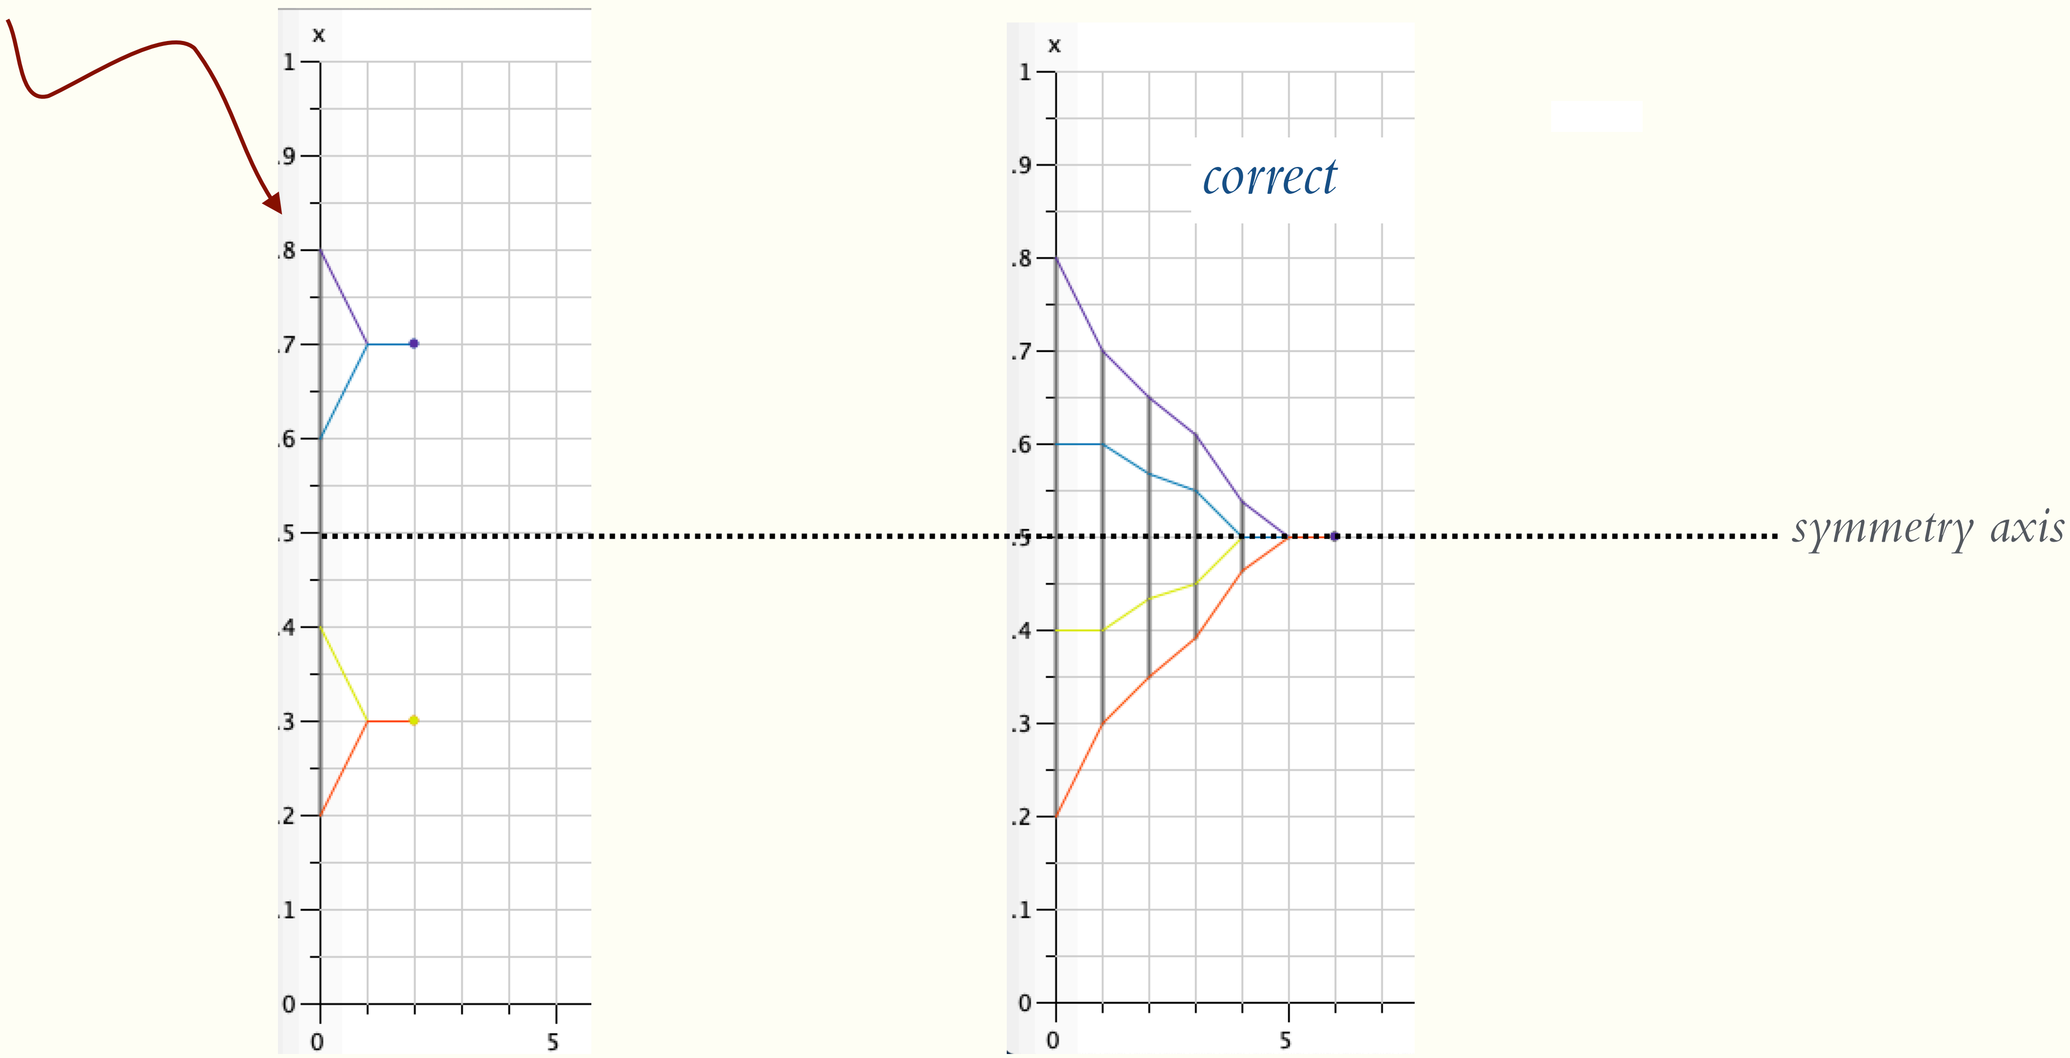

start profile: 0.2, 0.4, 0.6, 0.8
confidence level $\epsilon = 0.2$

Note: In the example here, the opinions are <u>exactly</u> on the borders of the confidence interval of other opinions. My search algorithm generates this numerical situations again and again.

Excursion: Floating-point arithmetic

Floating point arithmetic: The enemy within

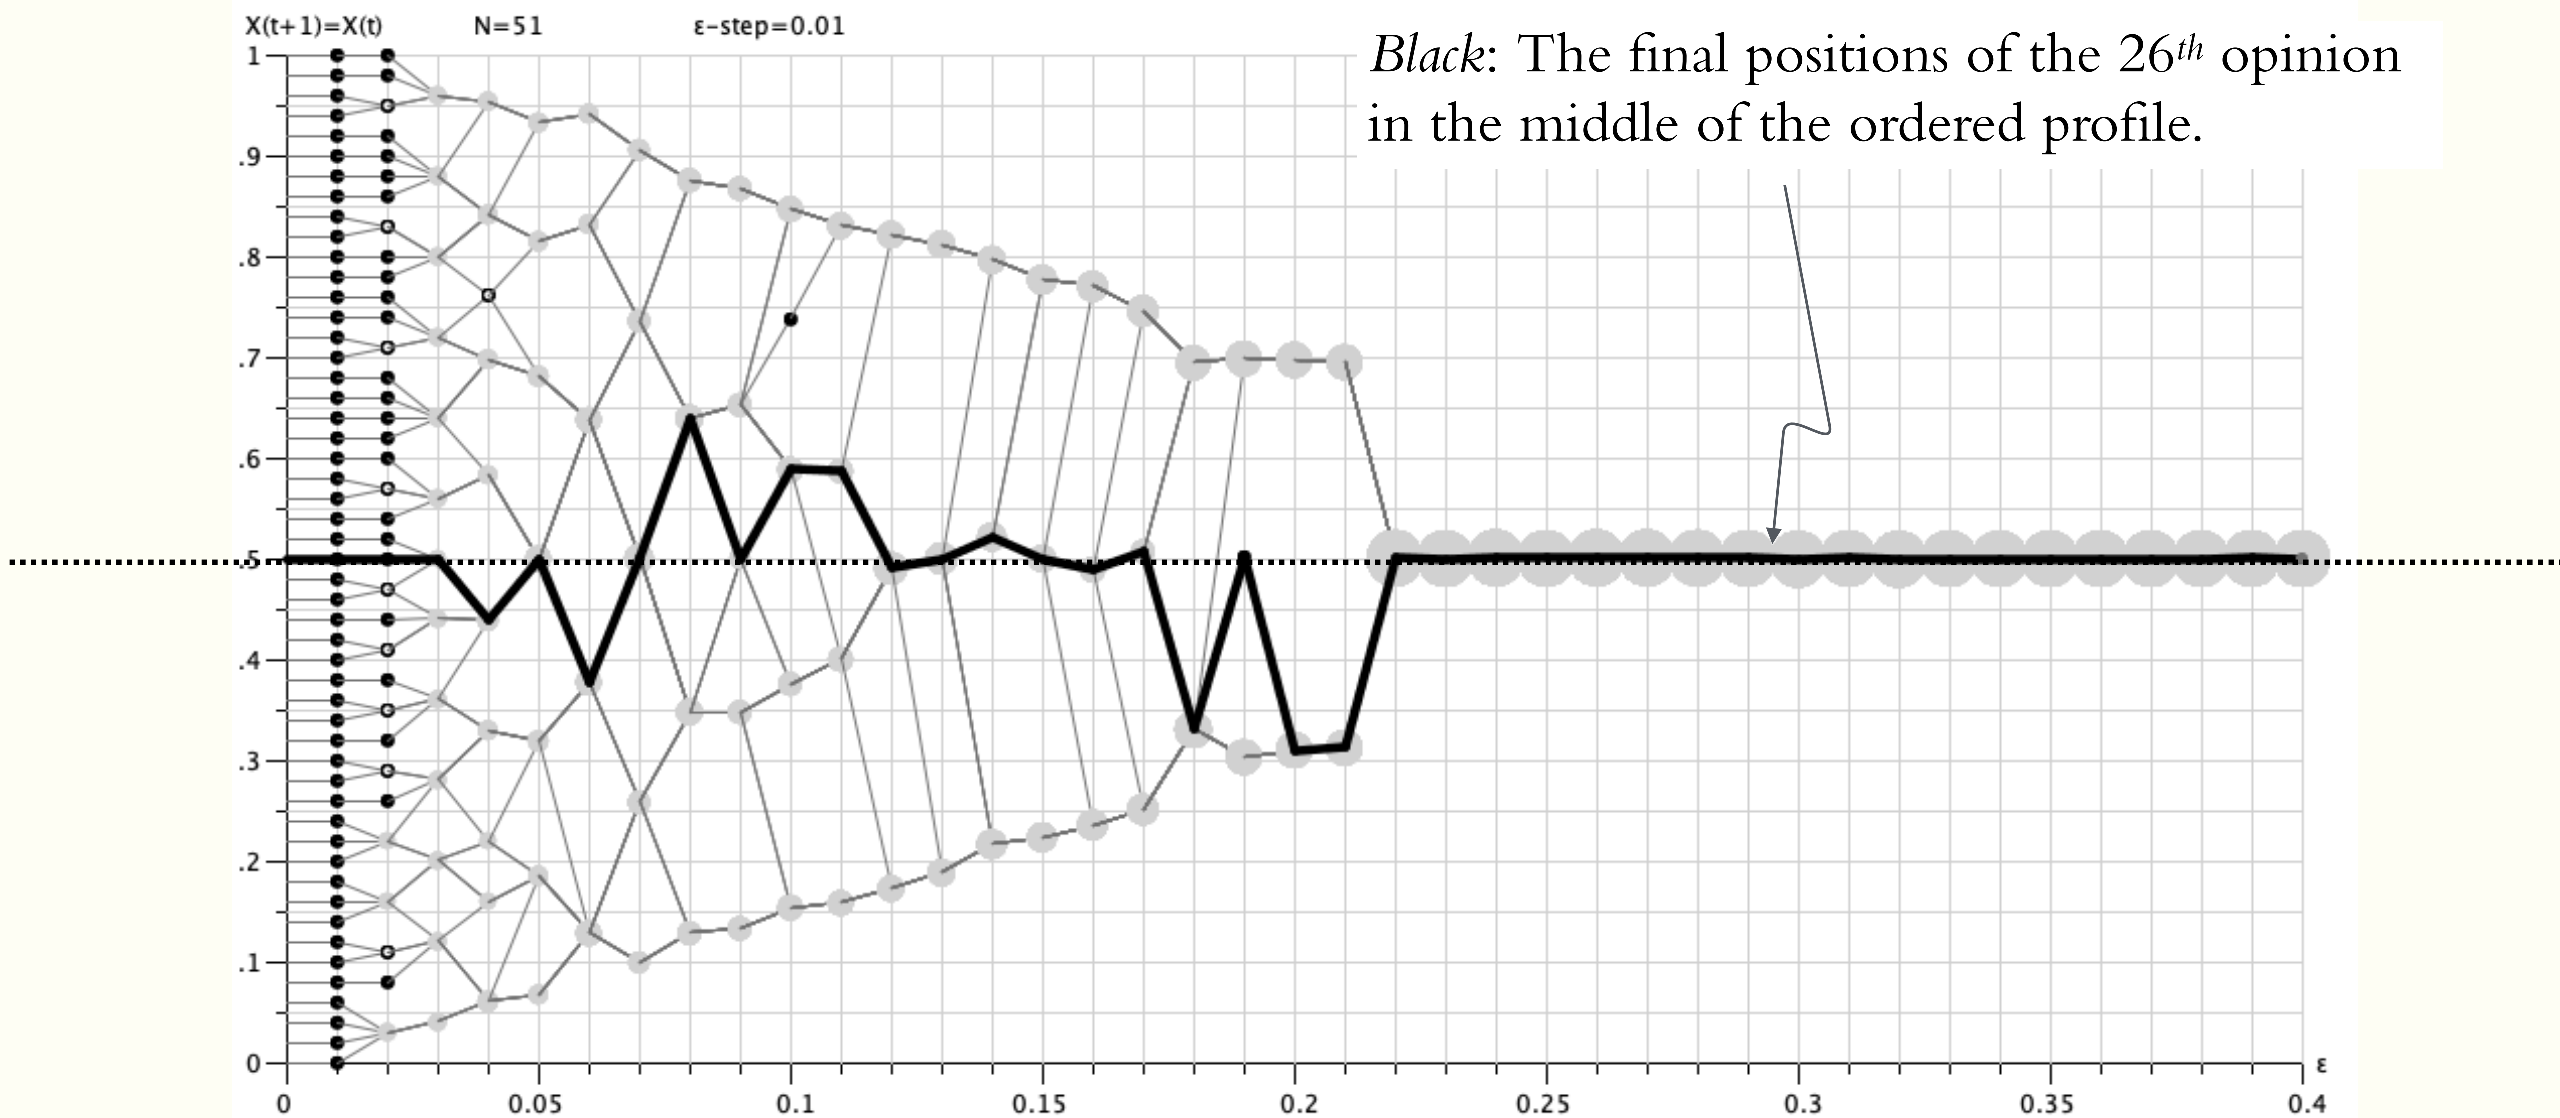

ε-*diagram* for a regular value start distribution with 51 opinions. The i^{th} opinion is $(i-1)/(n-1)= i / 50$. Since ε increases by ε = 0, 0.01, 0.02, …, 0.5, at the start lots of opinions are exactly at the bounds of confidence of other opinions.

Any deviation of the black line from 0.5 and any deviation from a mirror symmetry around 0.5 of the whole diagram, is a sufficient condition, that numerically something somewhere went wrong in the computation of the underlying unique single runs based upon a certain ε.

Excursion: Floating-point arithmetic

Decimal versus binary periodicity

non-periodic, ‚innocent' looking

decimal numbers	*number of their non-periodic binary counterparts*
0.1, 0.2, … , 0.9	1 out of 9: 0.5
0.01, 0.02, …, 0.99	3 out of 99: 0.25, 0.5, 0,75
0.001, 0.002, …, 0.999	7 out of 999: 0.125, 0.25, 0.5, 0.625, 0.75, 0.875
0.0001, 0.0002, …, 0.9999	15 out of 9999
0.00001, 0.00002, …, 0.99999	31 out of 99999

Almost all of these numbers are *binary periodic* and can't be *exactly* represented in a *binary* floating point arithmetic. They will be rounded. As a consequence, they are either a tiny bit too small or a tiny bit too big.

Our example involves the numbers 0.2, 0.4, 0.6, 0.8. *None* of them is binary exact.

Way out: Fractional arithmetic with integers of arbitrary length.

All possible BC-processes for our example $X(0)$

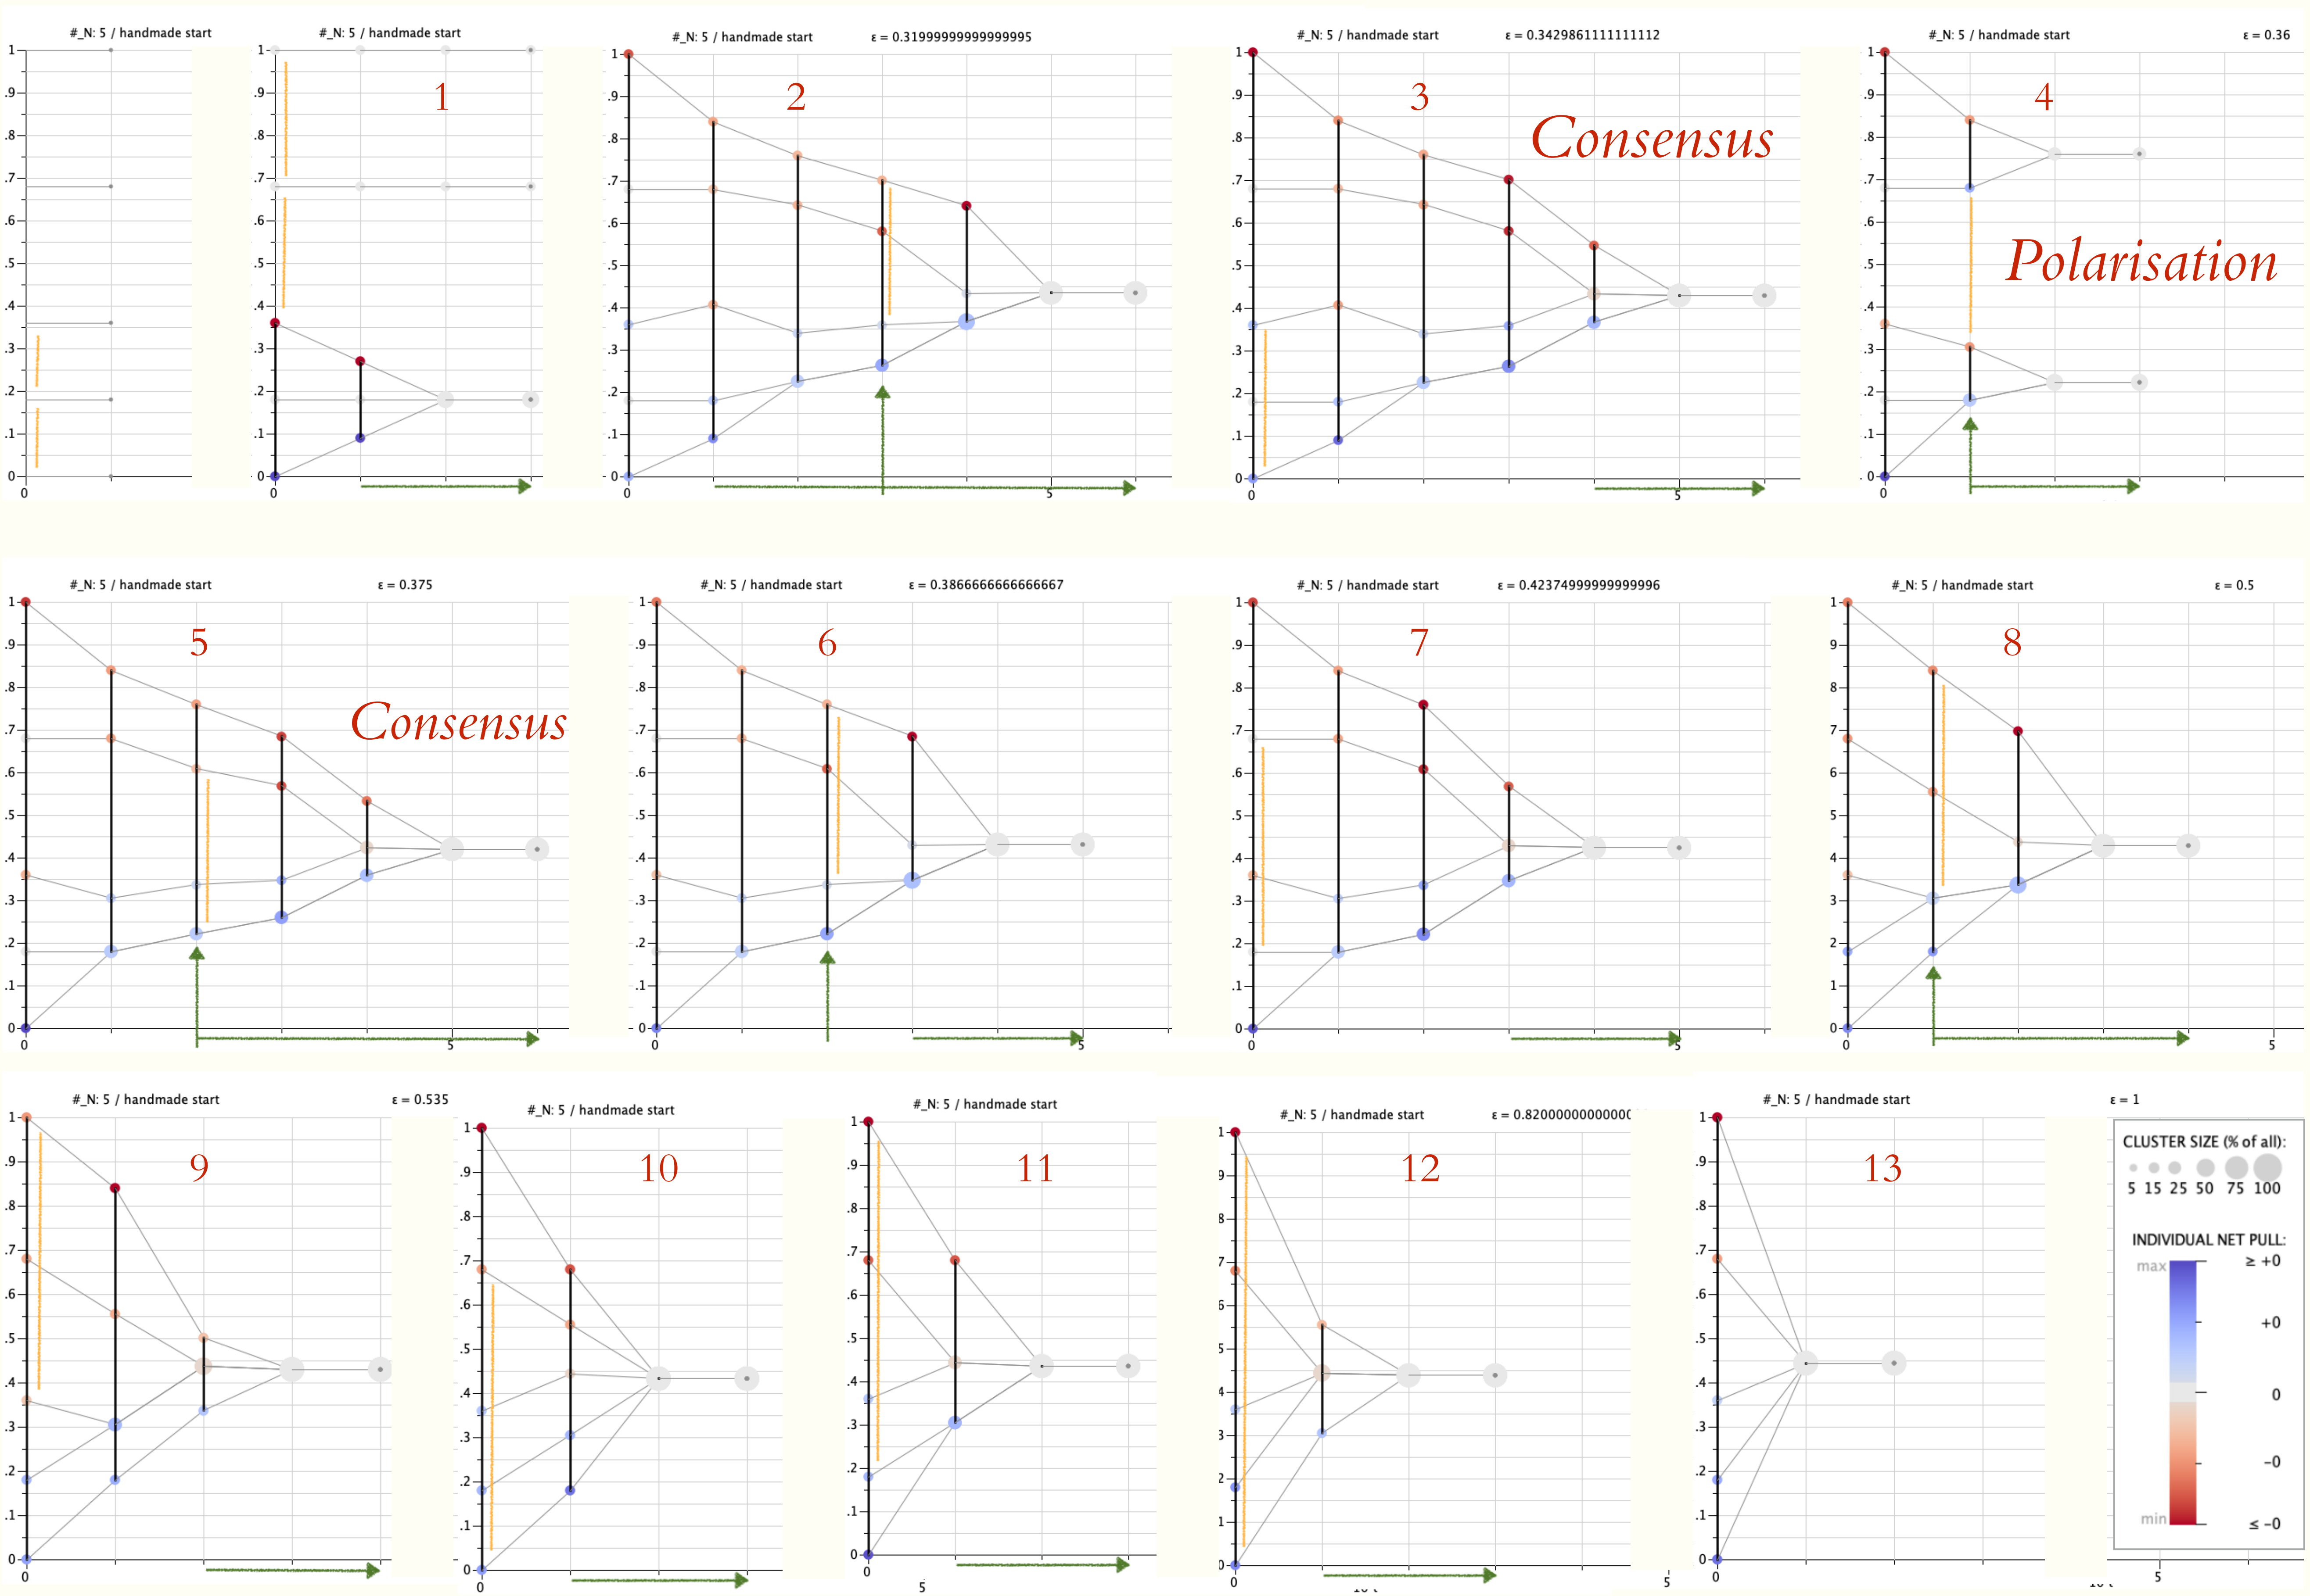

Discovery of non-monotonicities

JASSS

Jan Lorenz (2006)

Consensus Strikes Back in the Hegselmann-Krause Model of Continuous Opinion Dynamics Under Bounded Confidence

Journal of Artificial Societies and Social Simulation vol. 9, no. 1
<http://jasss.soc.surrey.ac.uk/9/1/8.html>

Computational results

1. For all start distributions $X(0)$ there exists only a *finite* number r of ϵ-values $0 < \epsilon_1^* < \epsilon_2^*, \ \ldots, \epsilon_{r-1}^* < \epsilon_r^* \leq 1$ that *make a difference* for the dynamics. I call them *ϵ-switches.*

2. The switches partition the unit interval [0,1] in a sequence of right-open intervals (except for the last one) $\langle\ [0, \epsilon_1^*)\ ,\ [\epsilon_1^*, \epsilon_2^*)\ ,\ \ldots,\ [\epsilon_r^*, 1]\ \rangle$. For all ϵ-values *within* such an interval, the resulting BC-dynamics *is the same*.

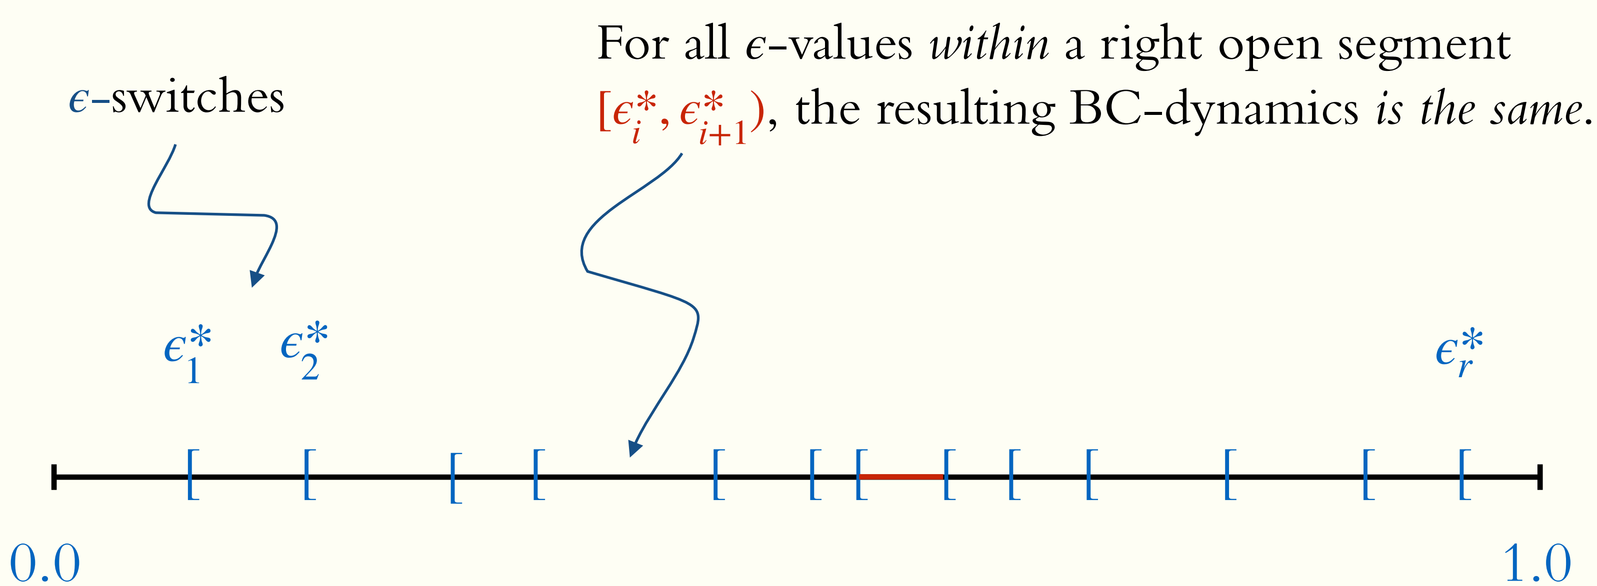

<u>Technical note:</u>
The algorithm that finds all switches requires an *exact* fractional arithmetic and support for integers of arbitrary length (BigInt). It is hopeless to try it with the usual floating-point arithmetic implemented in the *FPU* according *IEEE 754.*

§3

ϵ-Switches of equidistant start distributions with $n = 2, 3, \ldots, 50$ agents

Universal results for equidistant start profiles:
Total number of ϵ-switches $n = 2,\ldots,50$

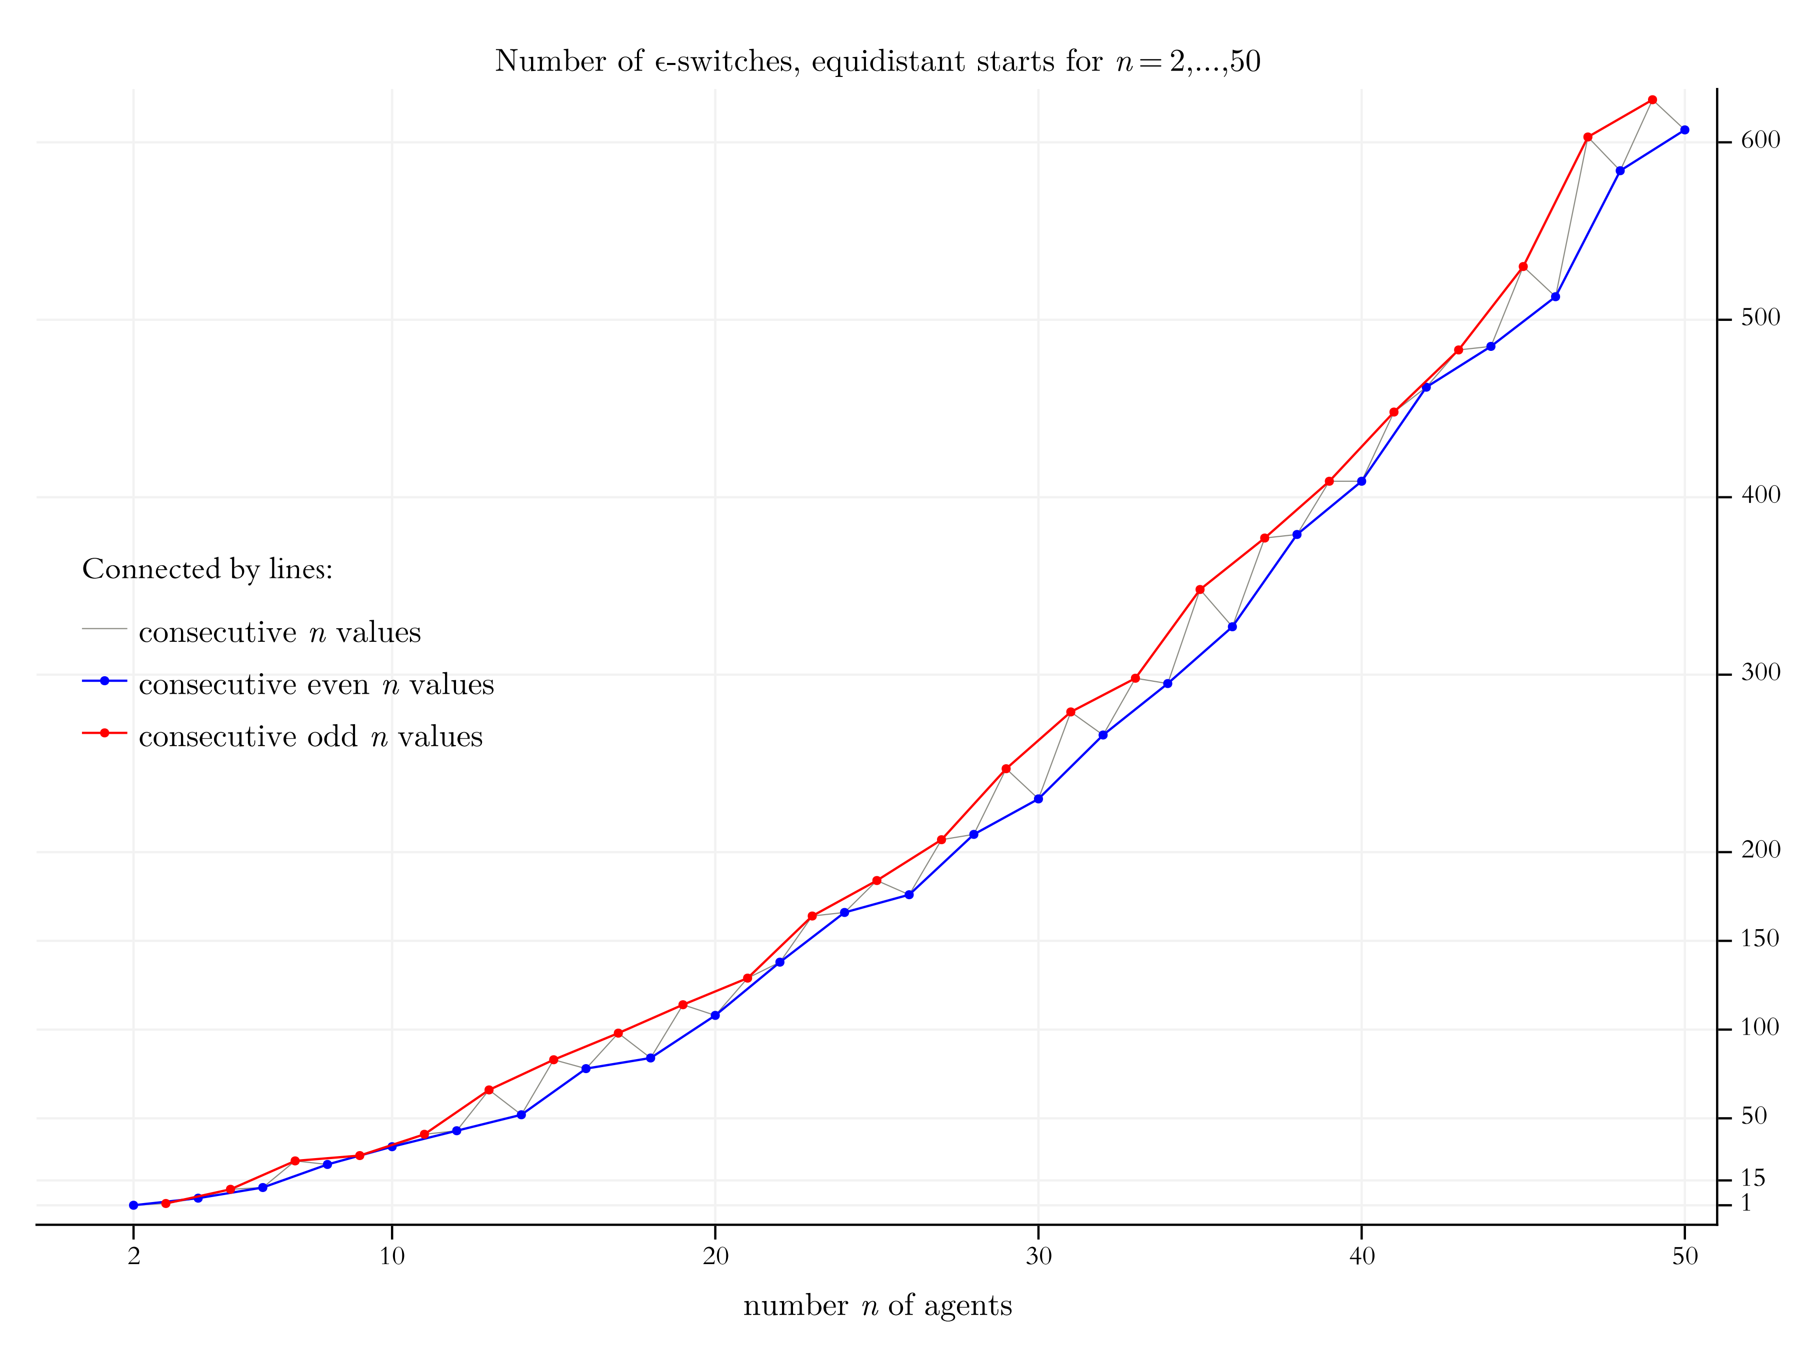

For increasing even values of n, and as well – but separately – for increasing odd values of n, the number of switches increases monotonically. In both cases the increase is more than linear. It looks like a polynomial increase. In most cases, but not always, the number of switches for an odd *n*, is greater than the number of switches for the even number (n + 1).

Universal results for equidistant start profiles:
Number of ϵ-switches from consent to 2 clusters

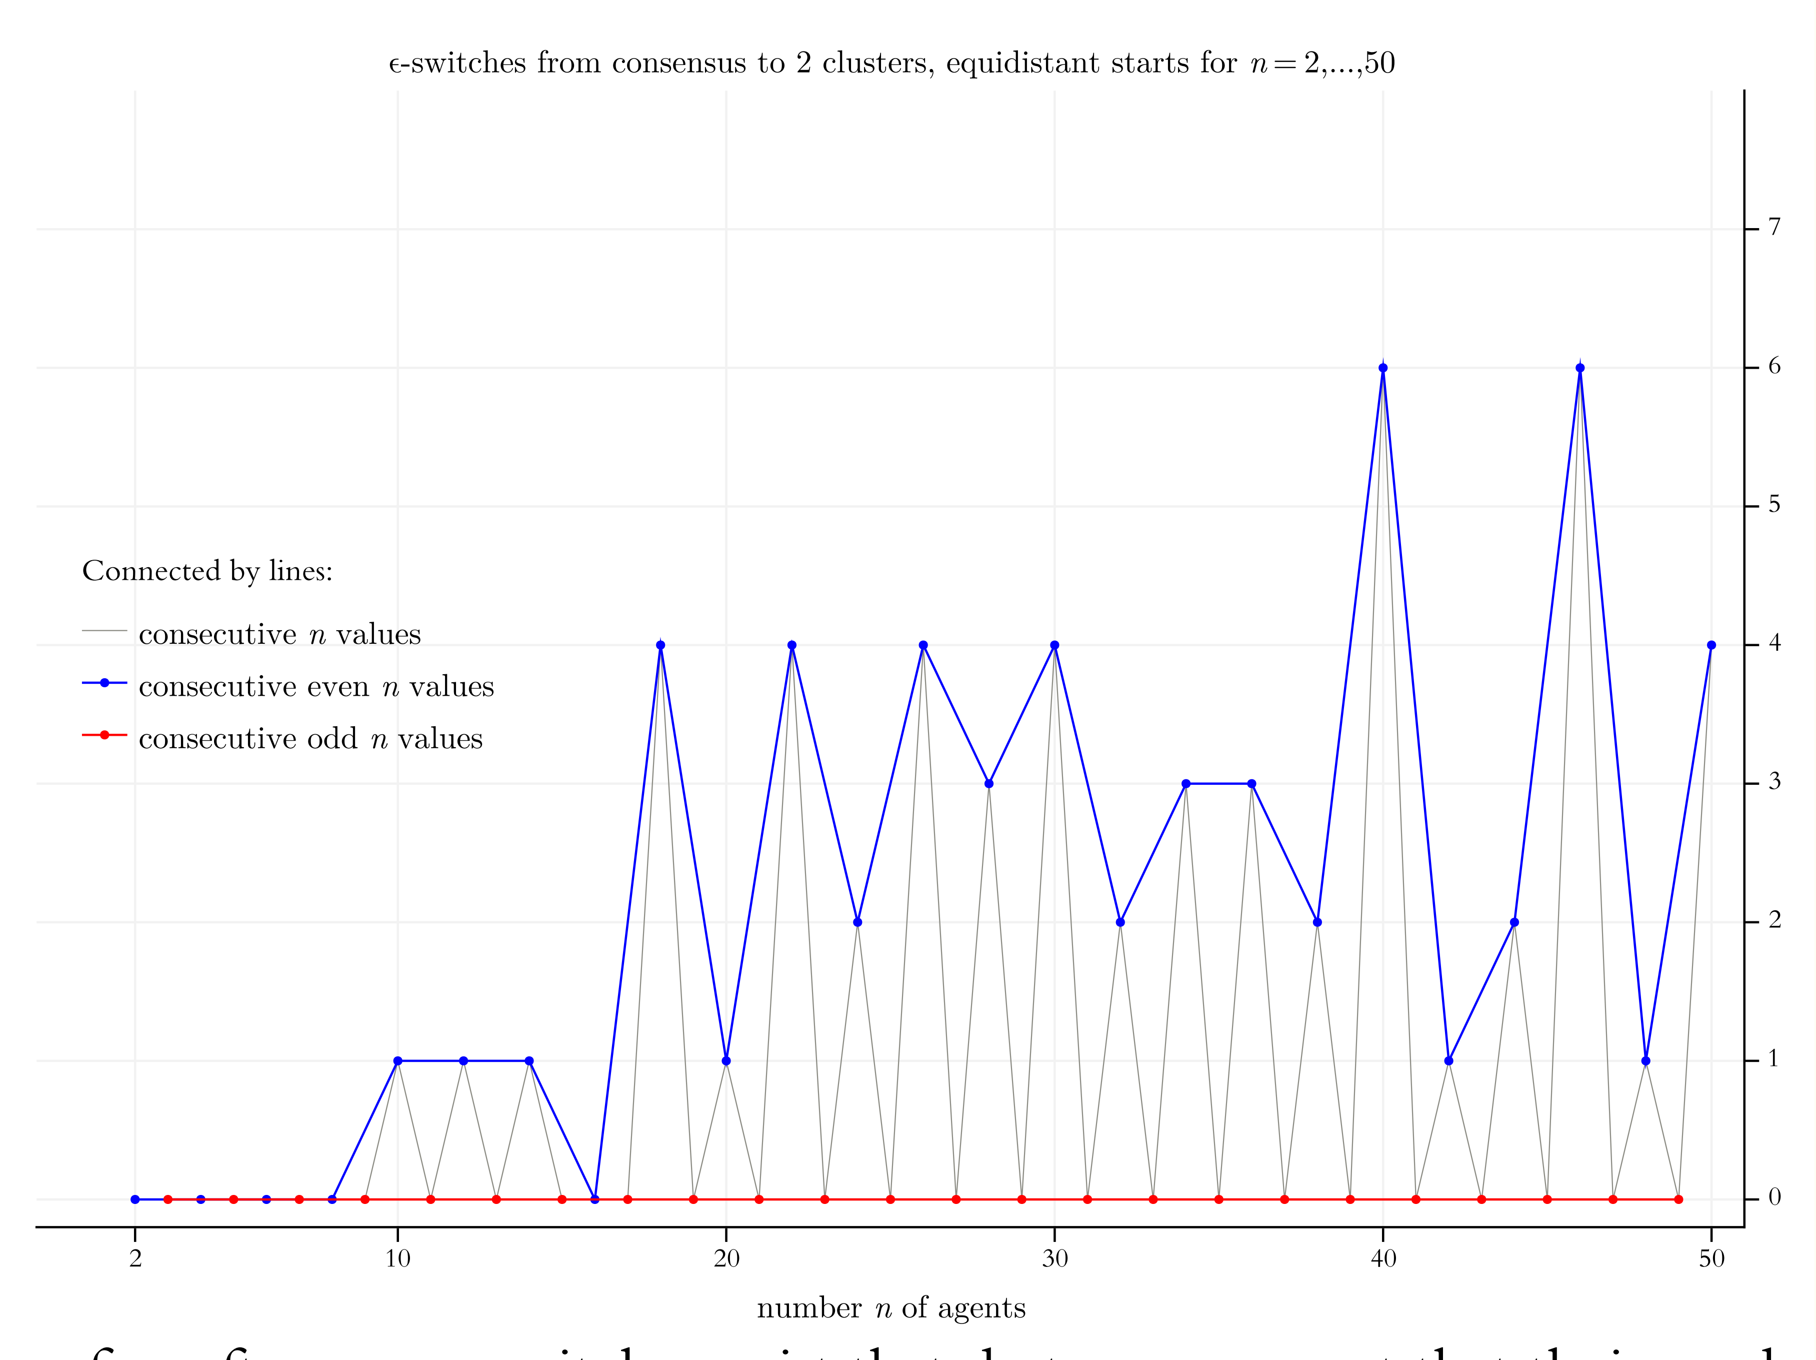

For even values of n, often many switches exist that destroy a consent that their predecessor switch generated. As even values of n become larger, one also seems to encounter larger numbers of such cases. For odd values of n, there are no switches that destroy a consensus and, at the same time, lead to polarisation in the strict sense of just two final clusters.

Universal results for equidistant start profiles:
Number of ϵ-switches from consent to 3 clusters

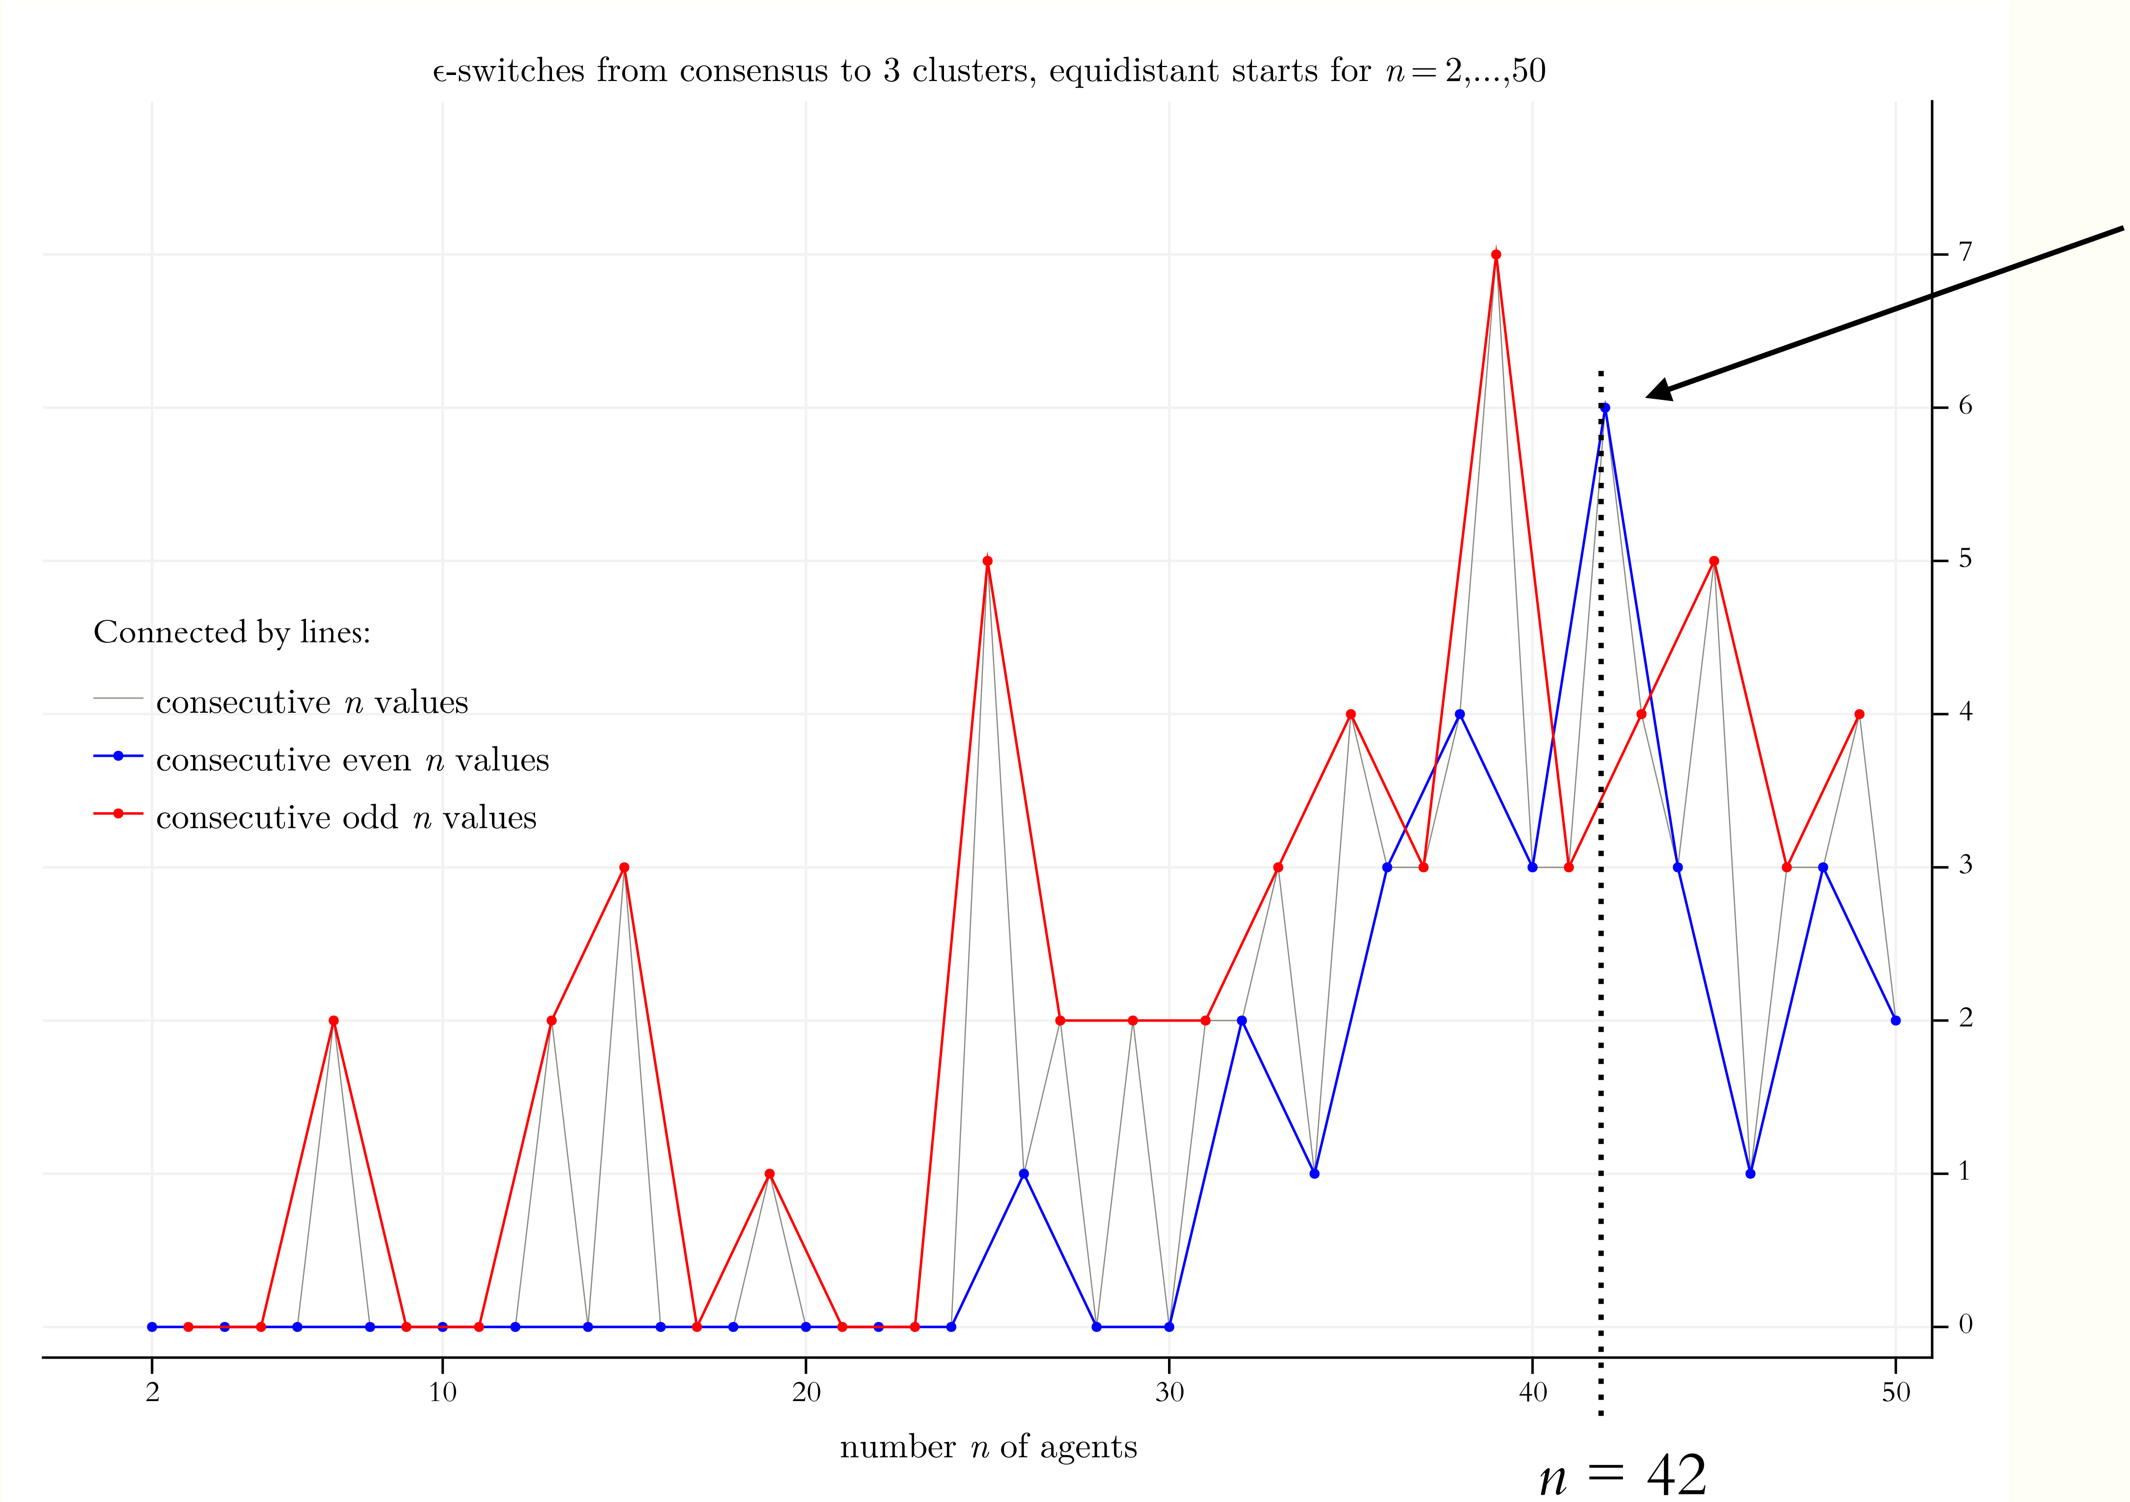

Special analysis will follow soon.

Universal results for equidistant start profiles:

Number of ϵ-switches with more final clusters

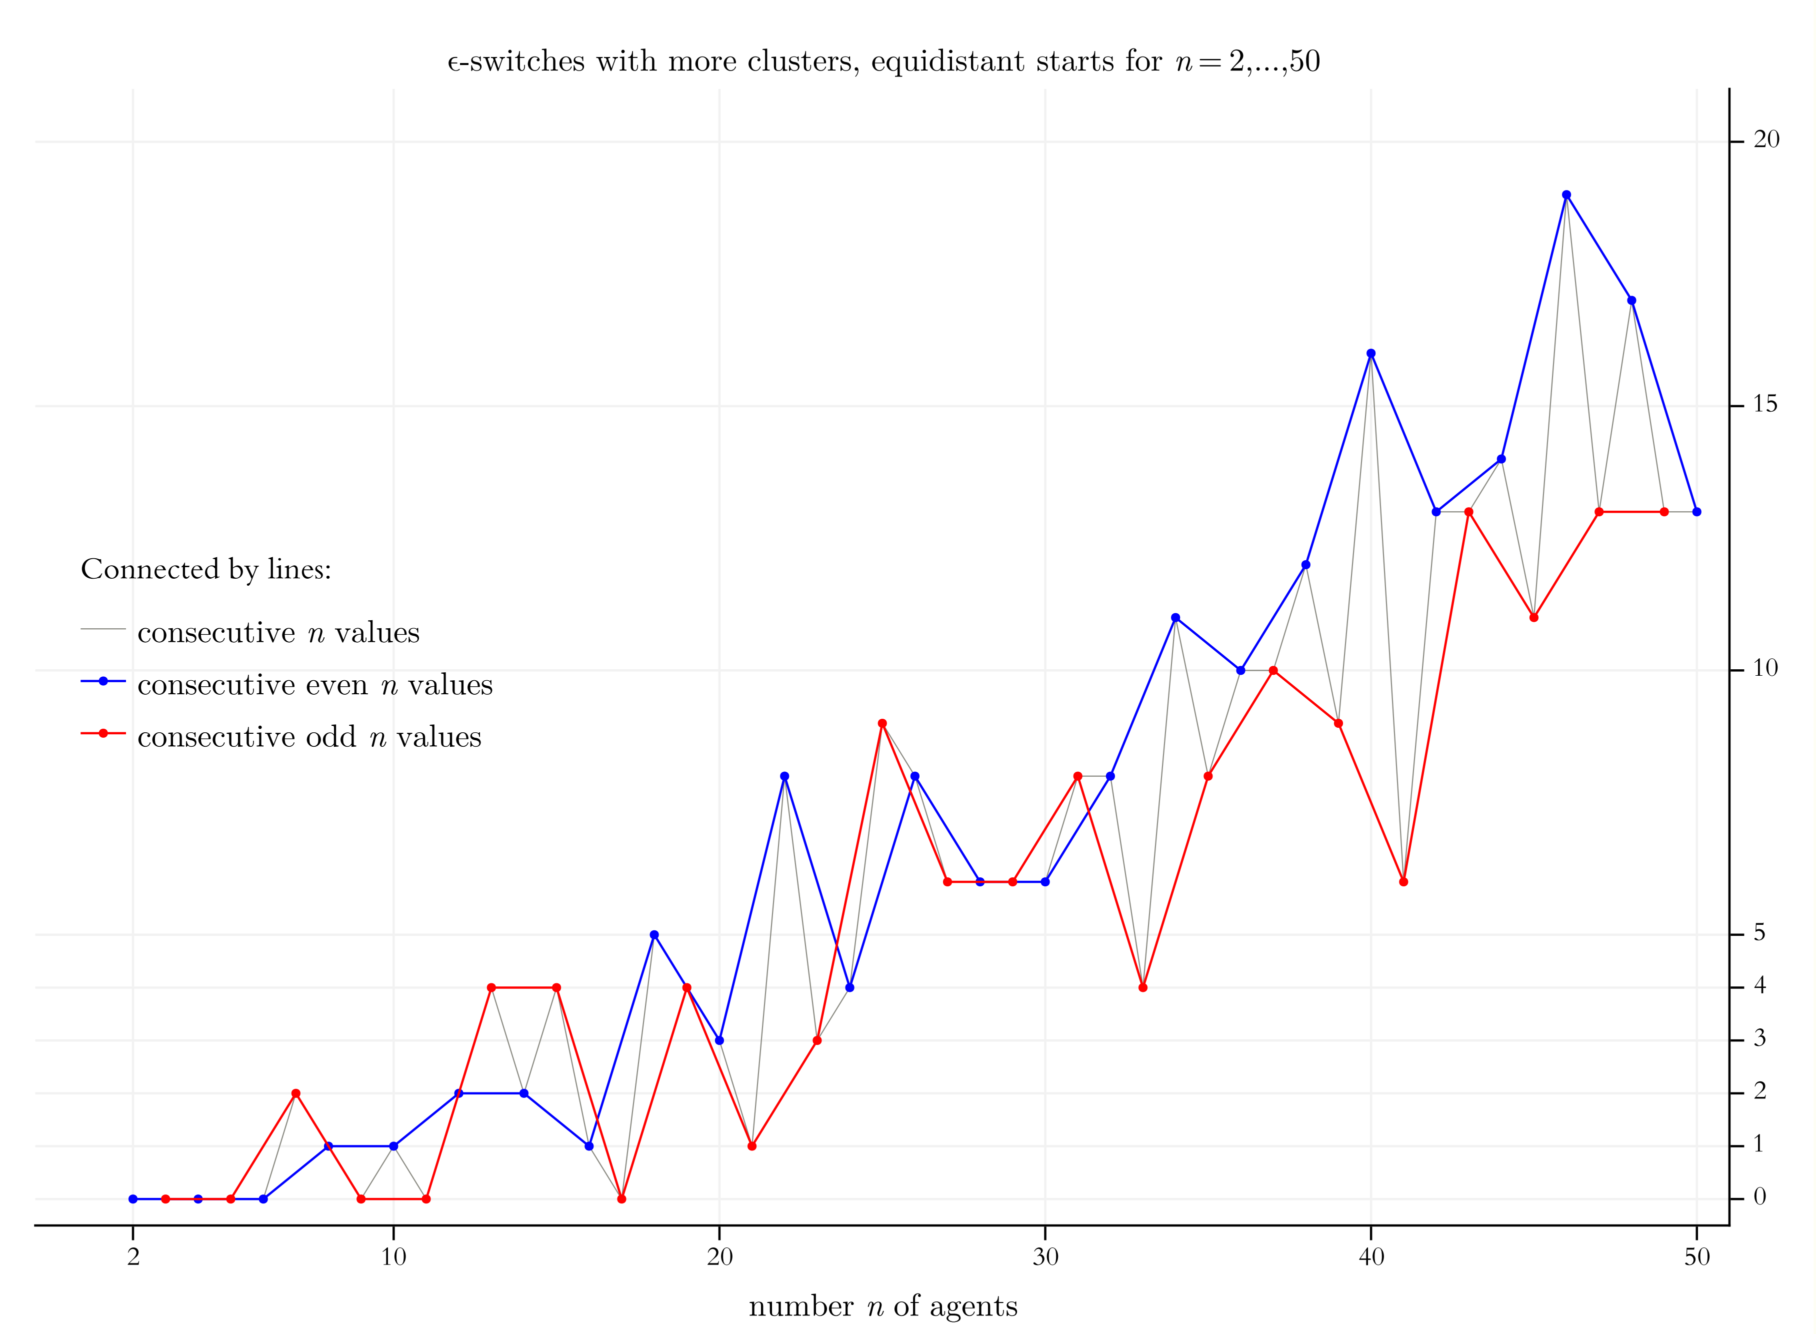

For increasing even and odd values of n, there is a non-monotonic tendency to occurrences of ever greater numbers of switches that, compared to their immediate predecessor, lead to more final clusters. In most, but not all cases, regular start distribution with an even value n have more such switches than the start distribution for the odd value $(n-1)$.

Universal results for equidistant start profiles:
Number of ϵ-switches with greater final profile width

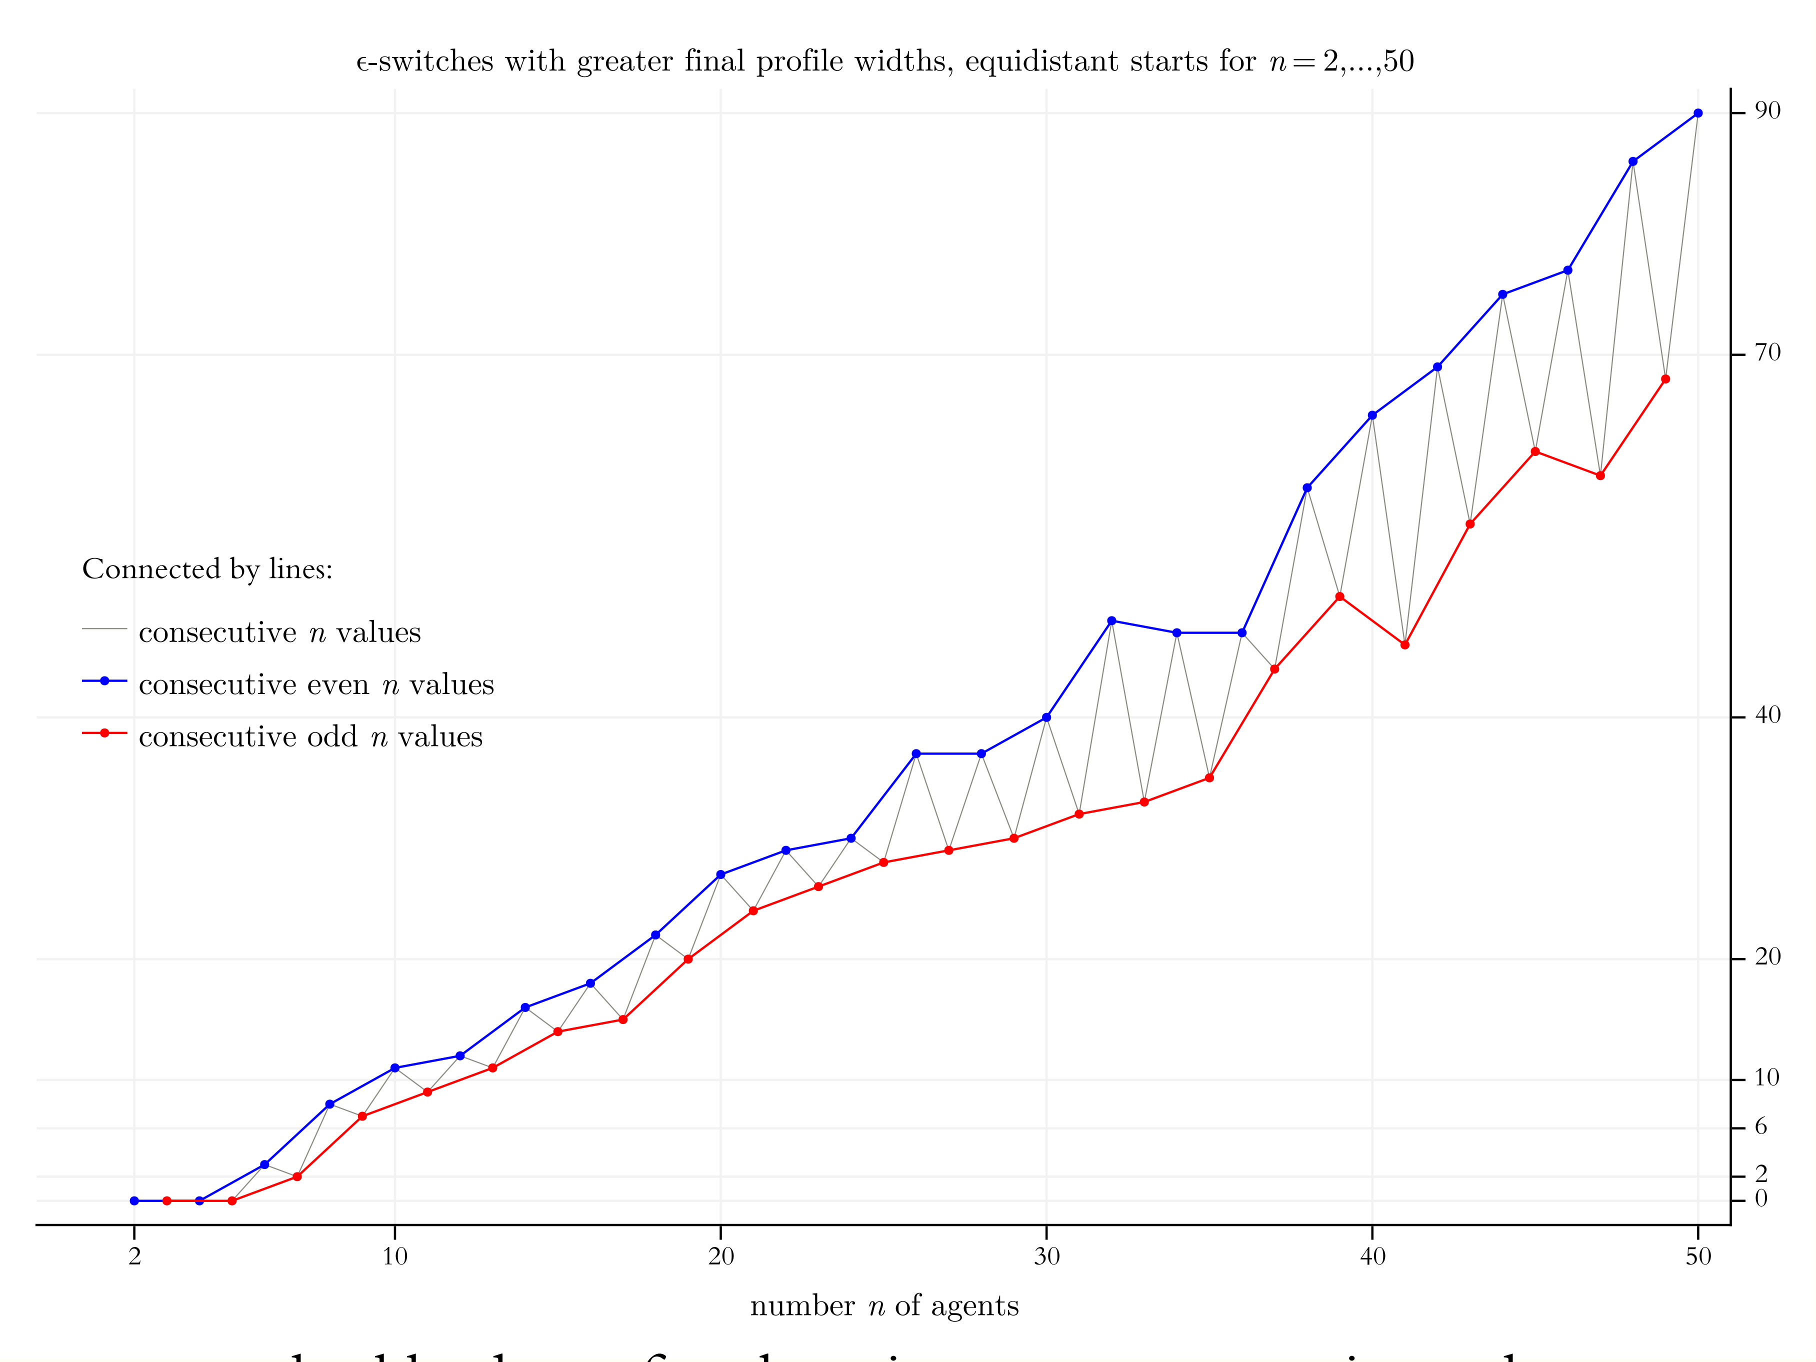

For increasing even and odd values of n, there is a non-monotonic tendency to ever greater numbers of switches that, compared to their immediate predecessor, lead to a larger final profile width. Except for very small values of n, regular start distribution with an even value n have more such switches than the start distribution for the odd value $(n-1)$.

§4
ϵ-Switch diagrams

Switch diagram: equidistant start, $n = 42$

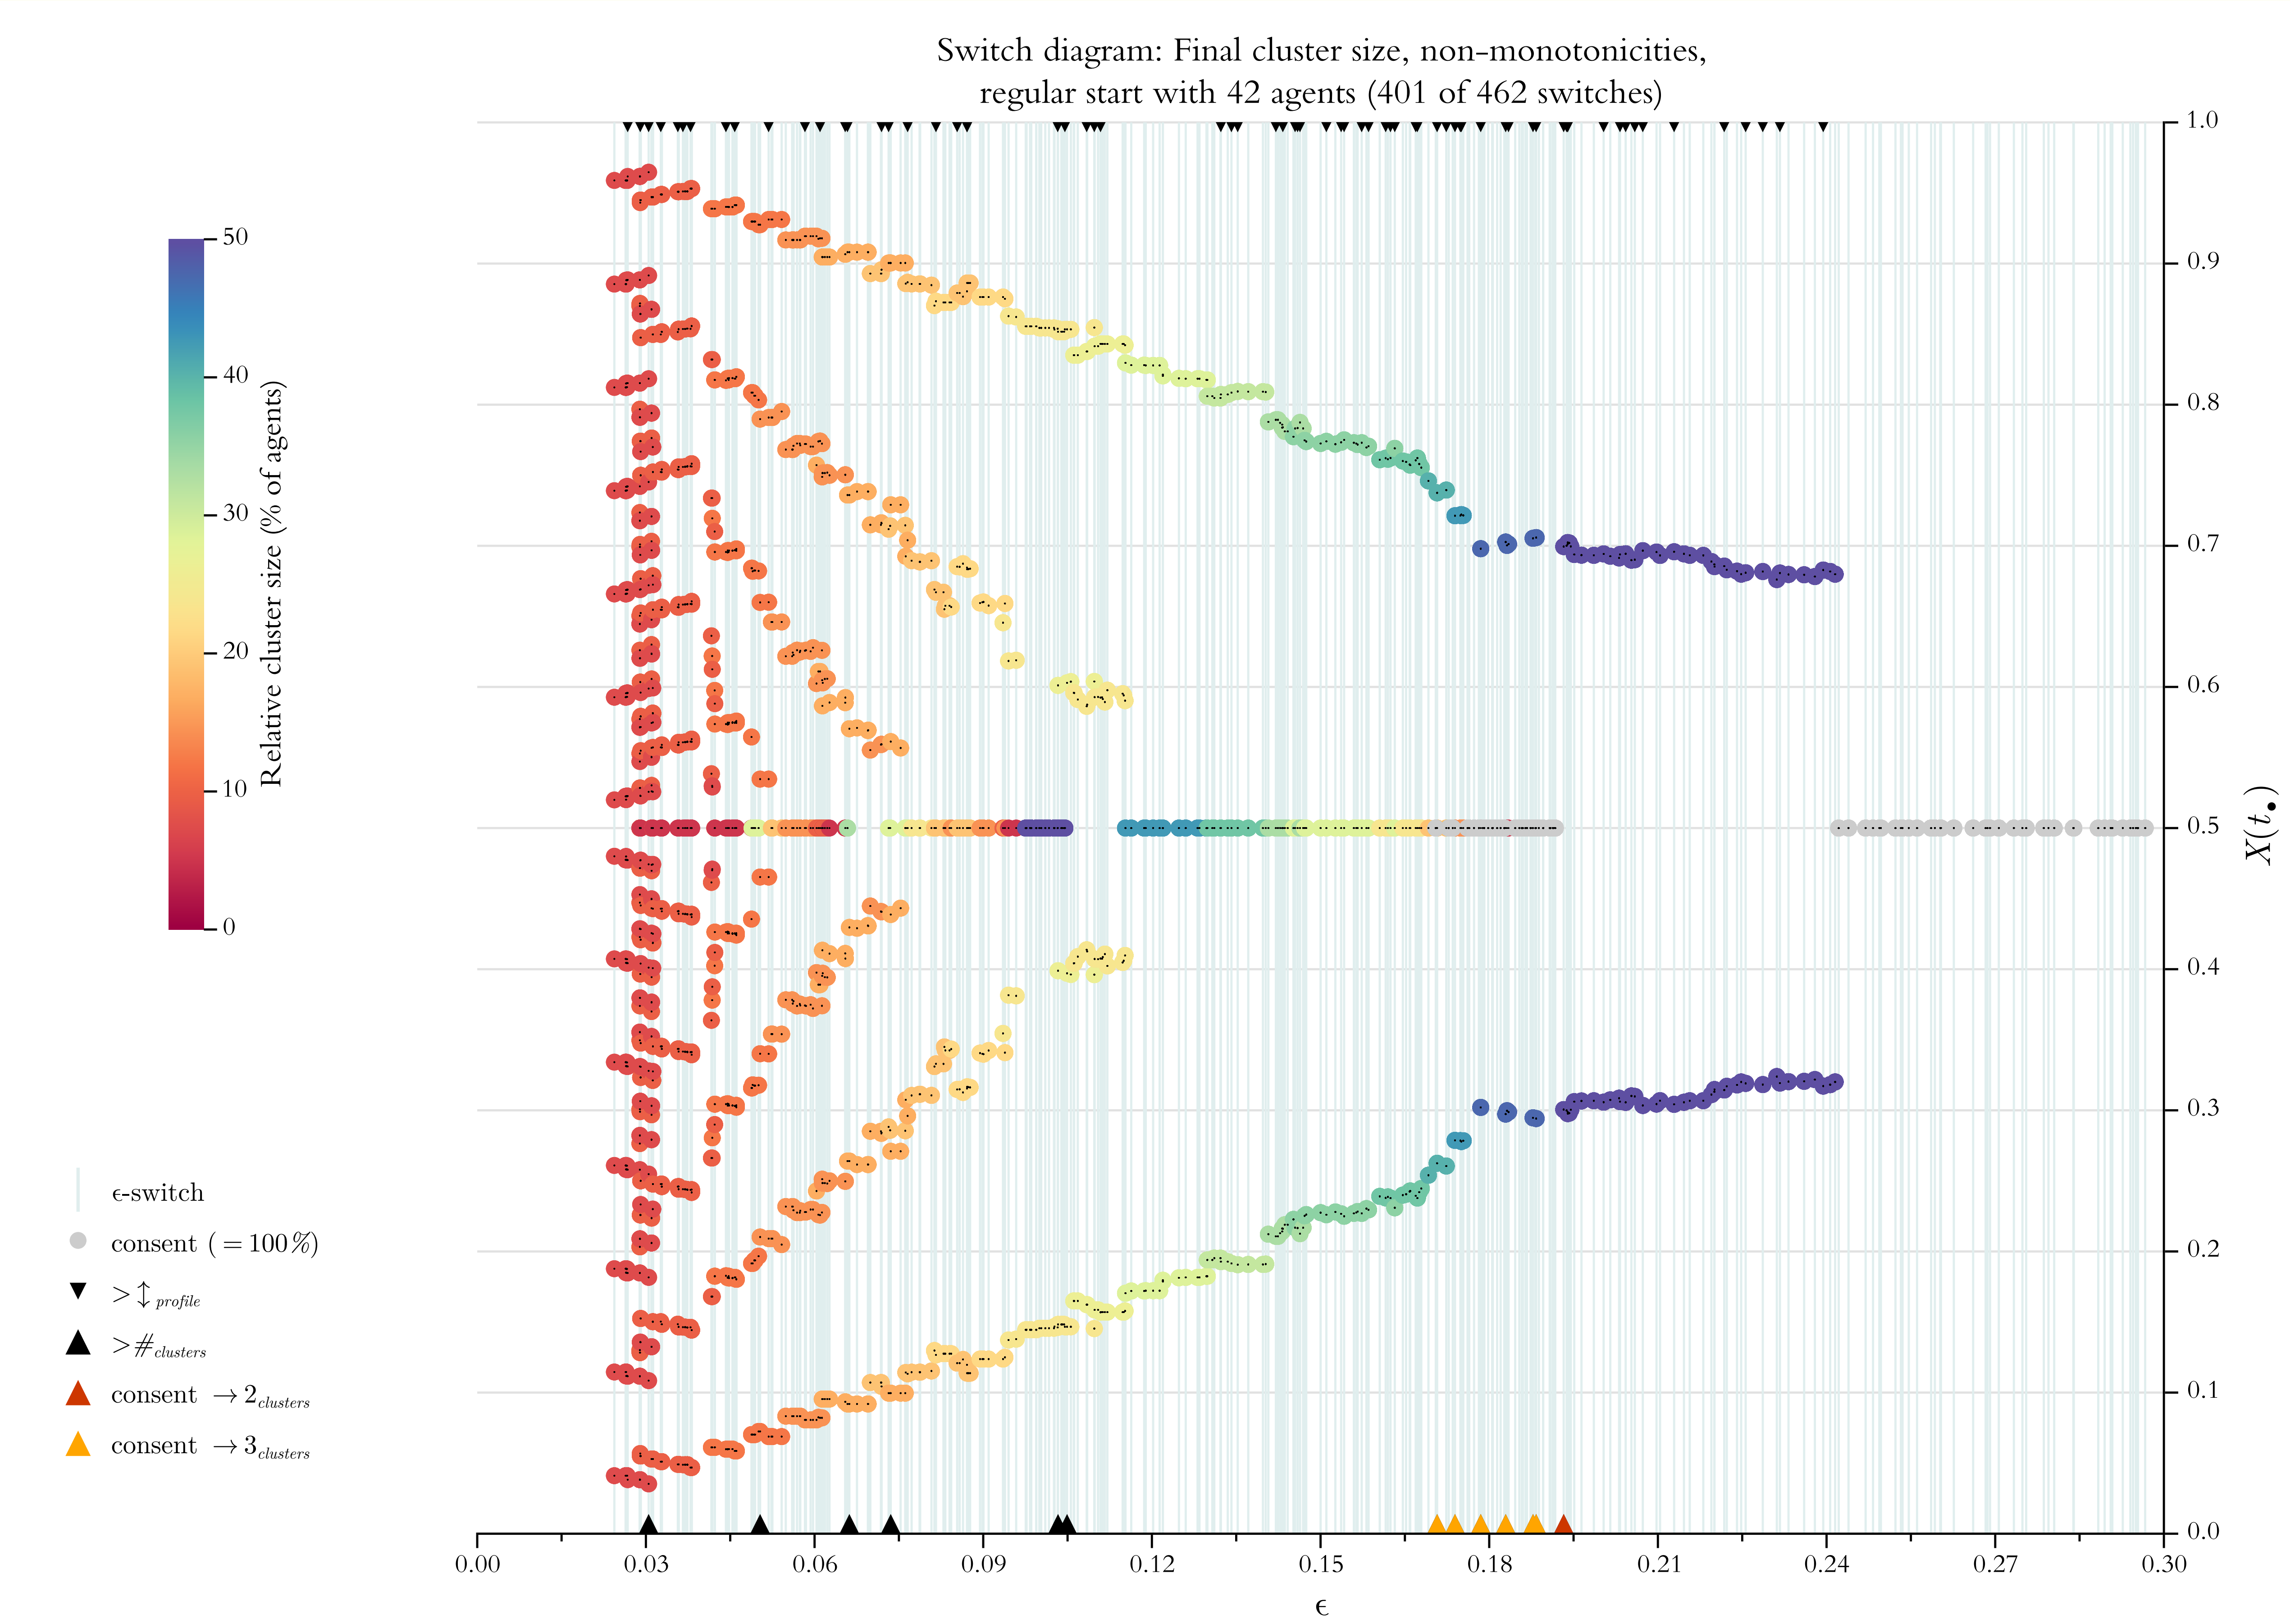

Switch diagram: equidistant start, $n = 42$

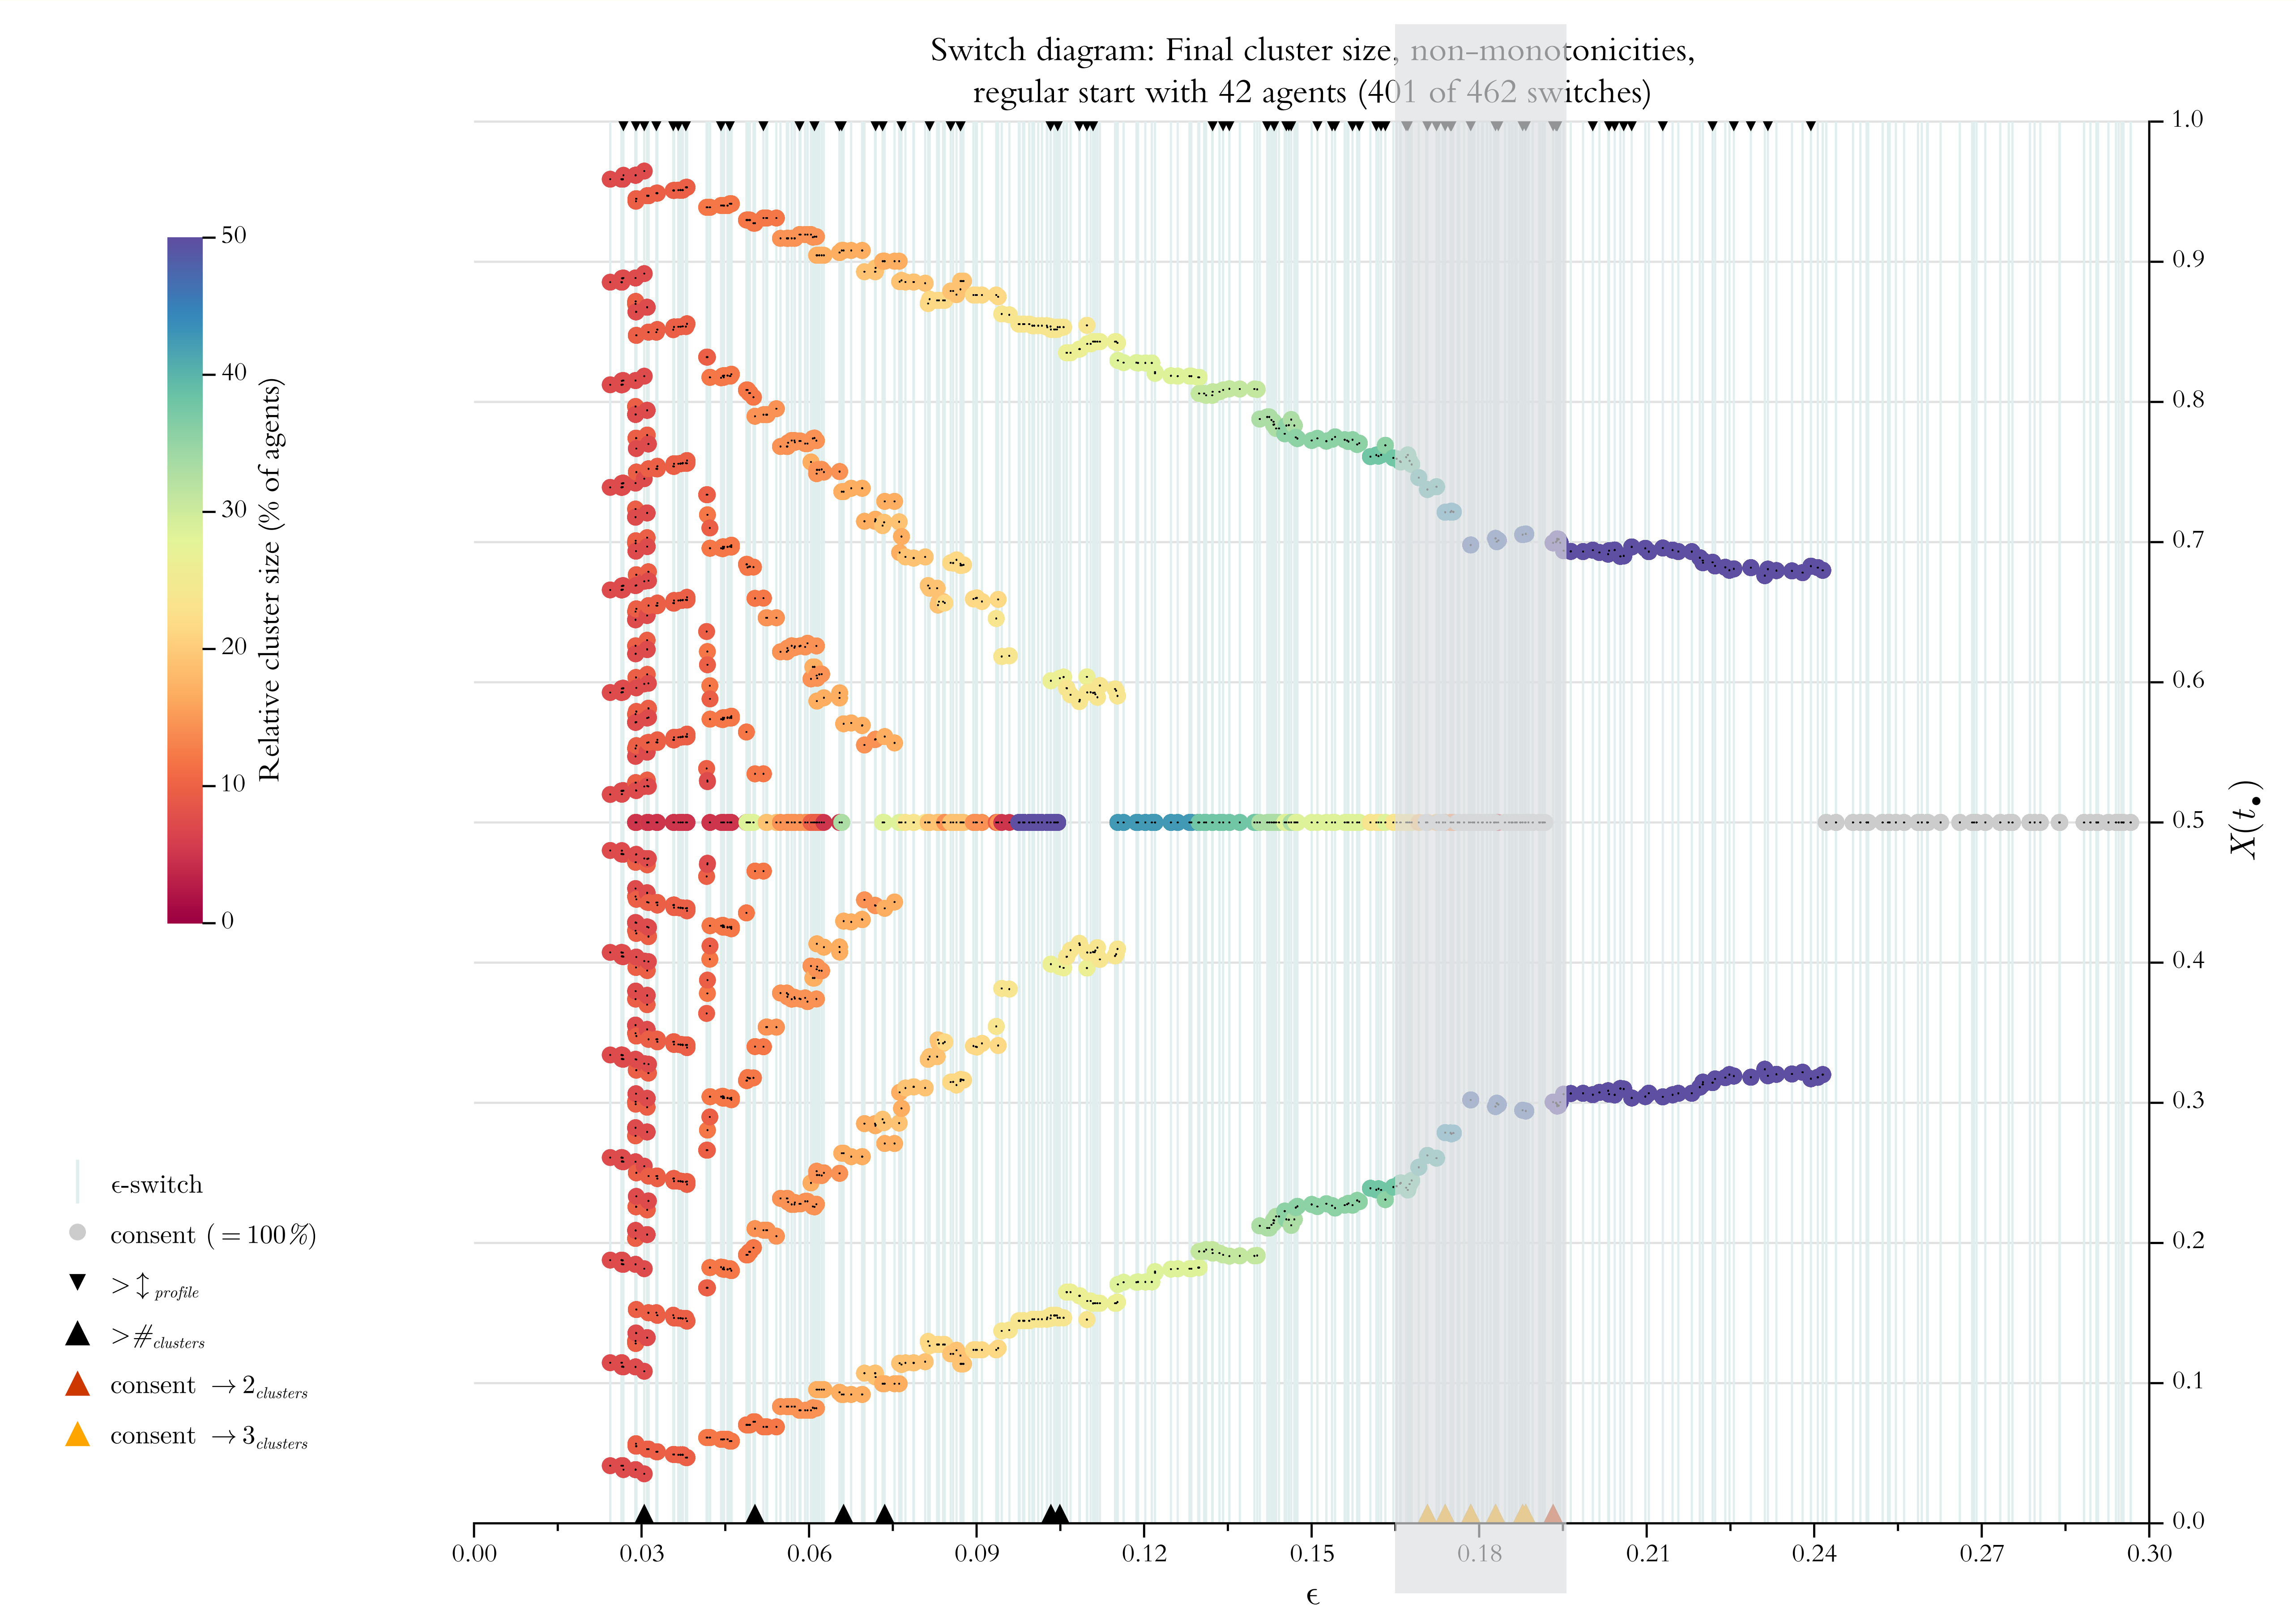

Switch diagram: equidistant start, $n = 42$

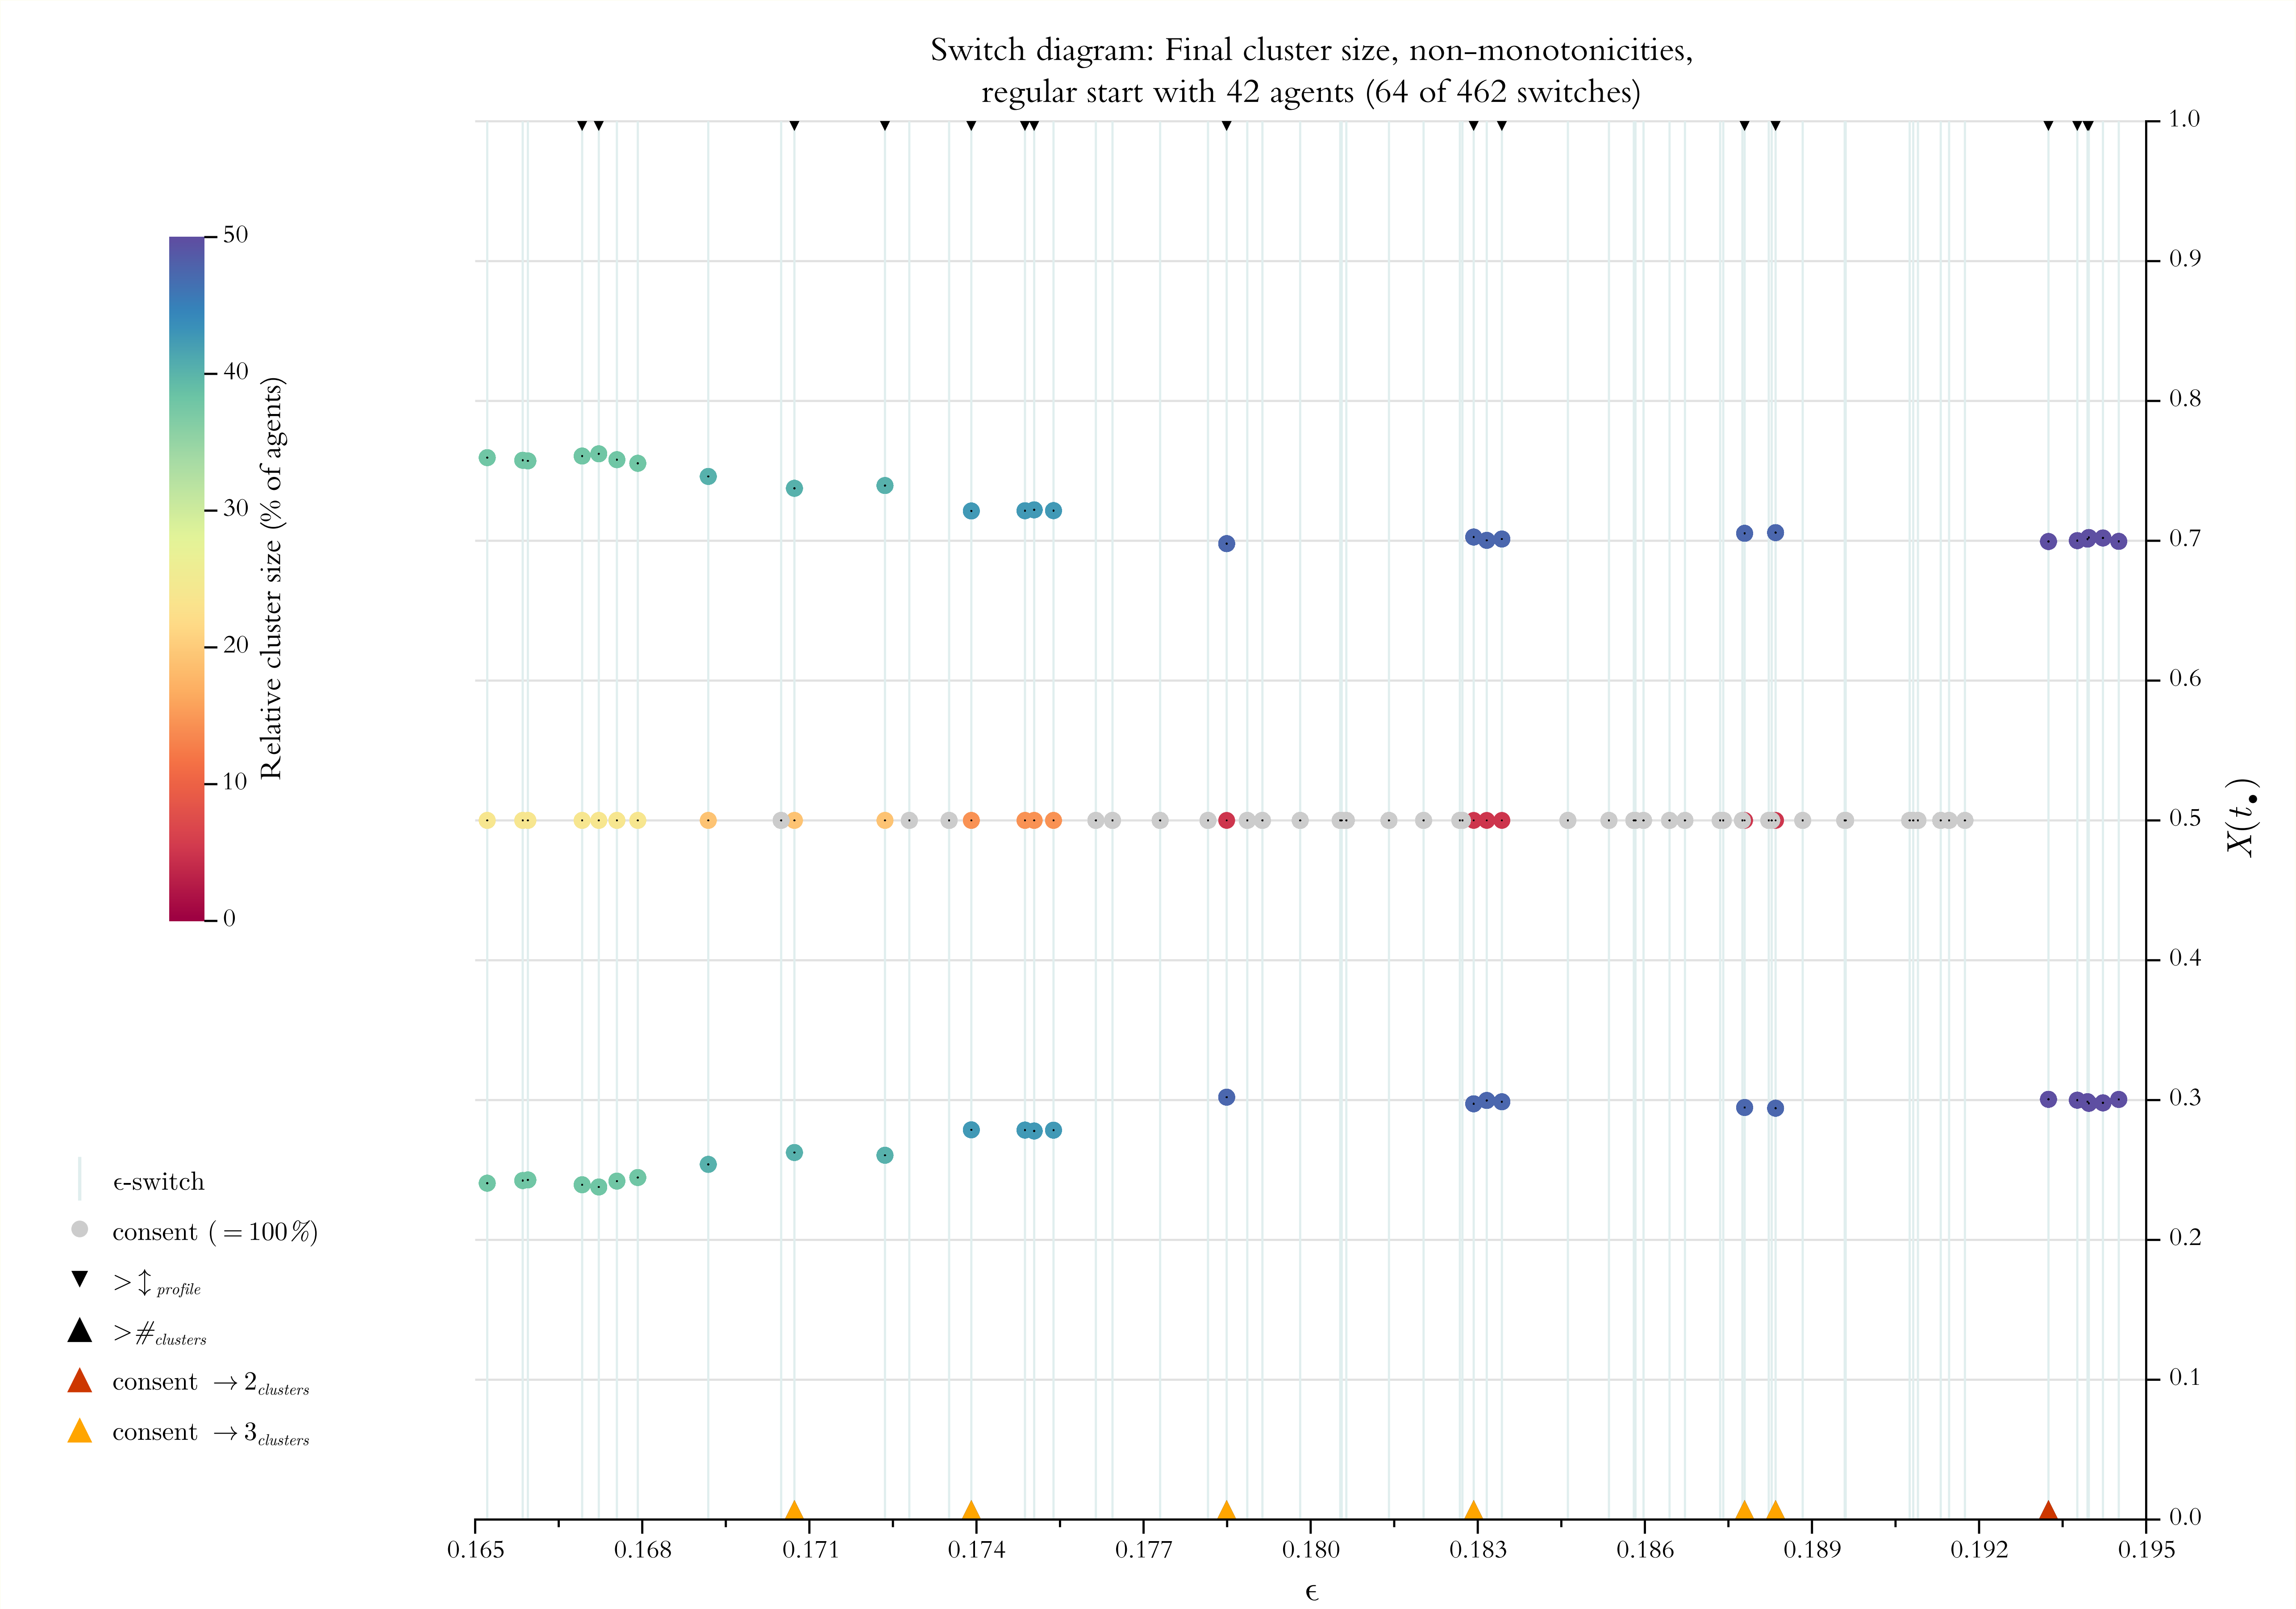

ϵ-*switch diagrams*: regular start, $n = 2,4,\ldots,80$

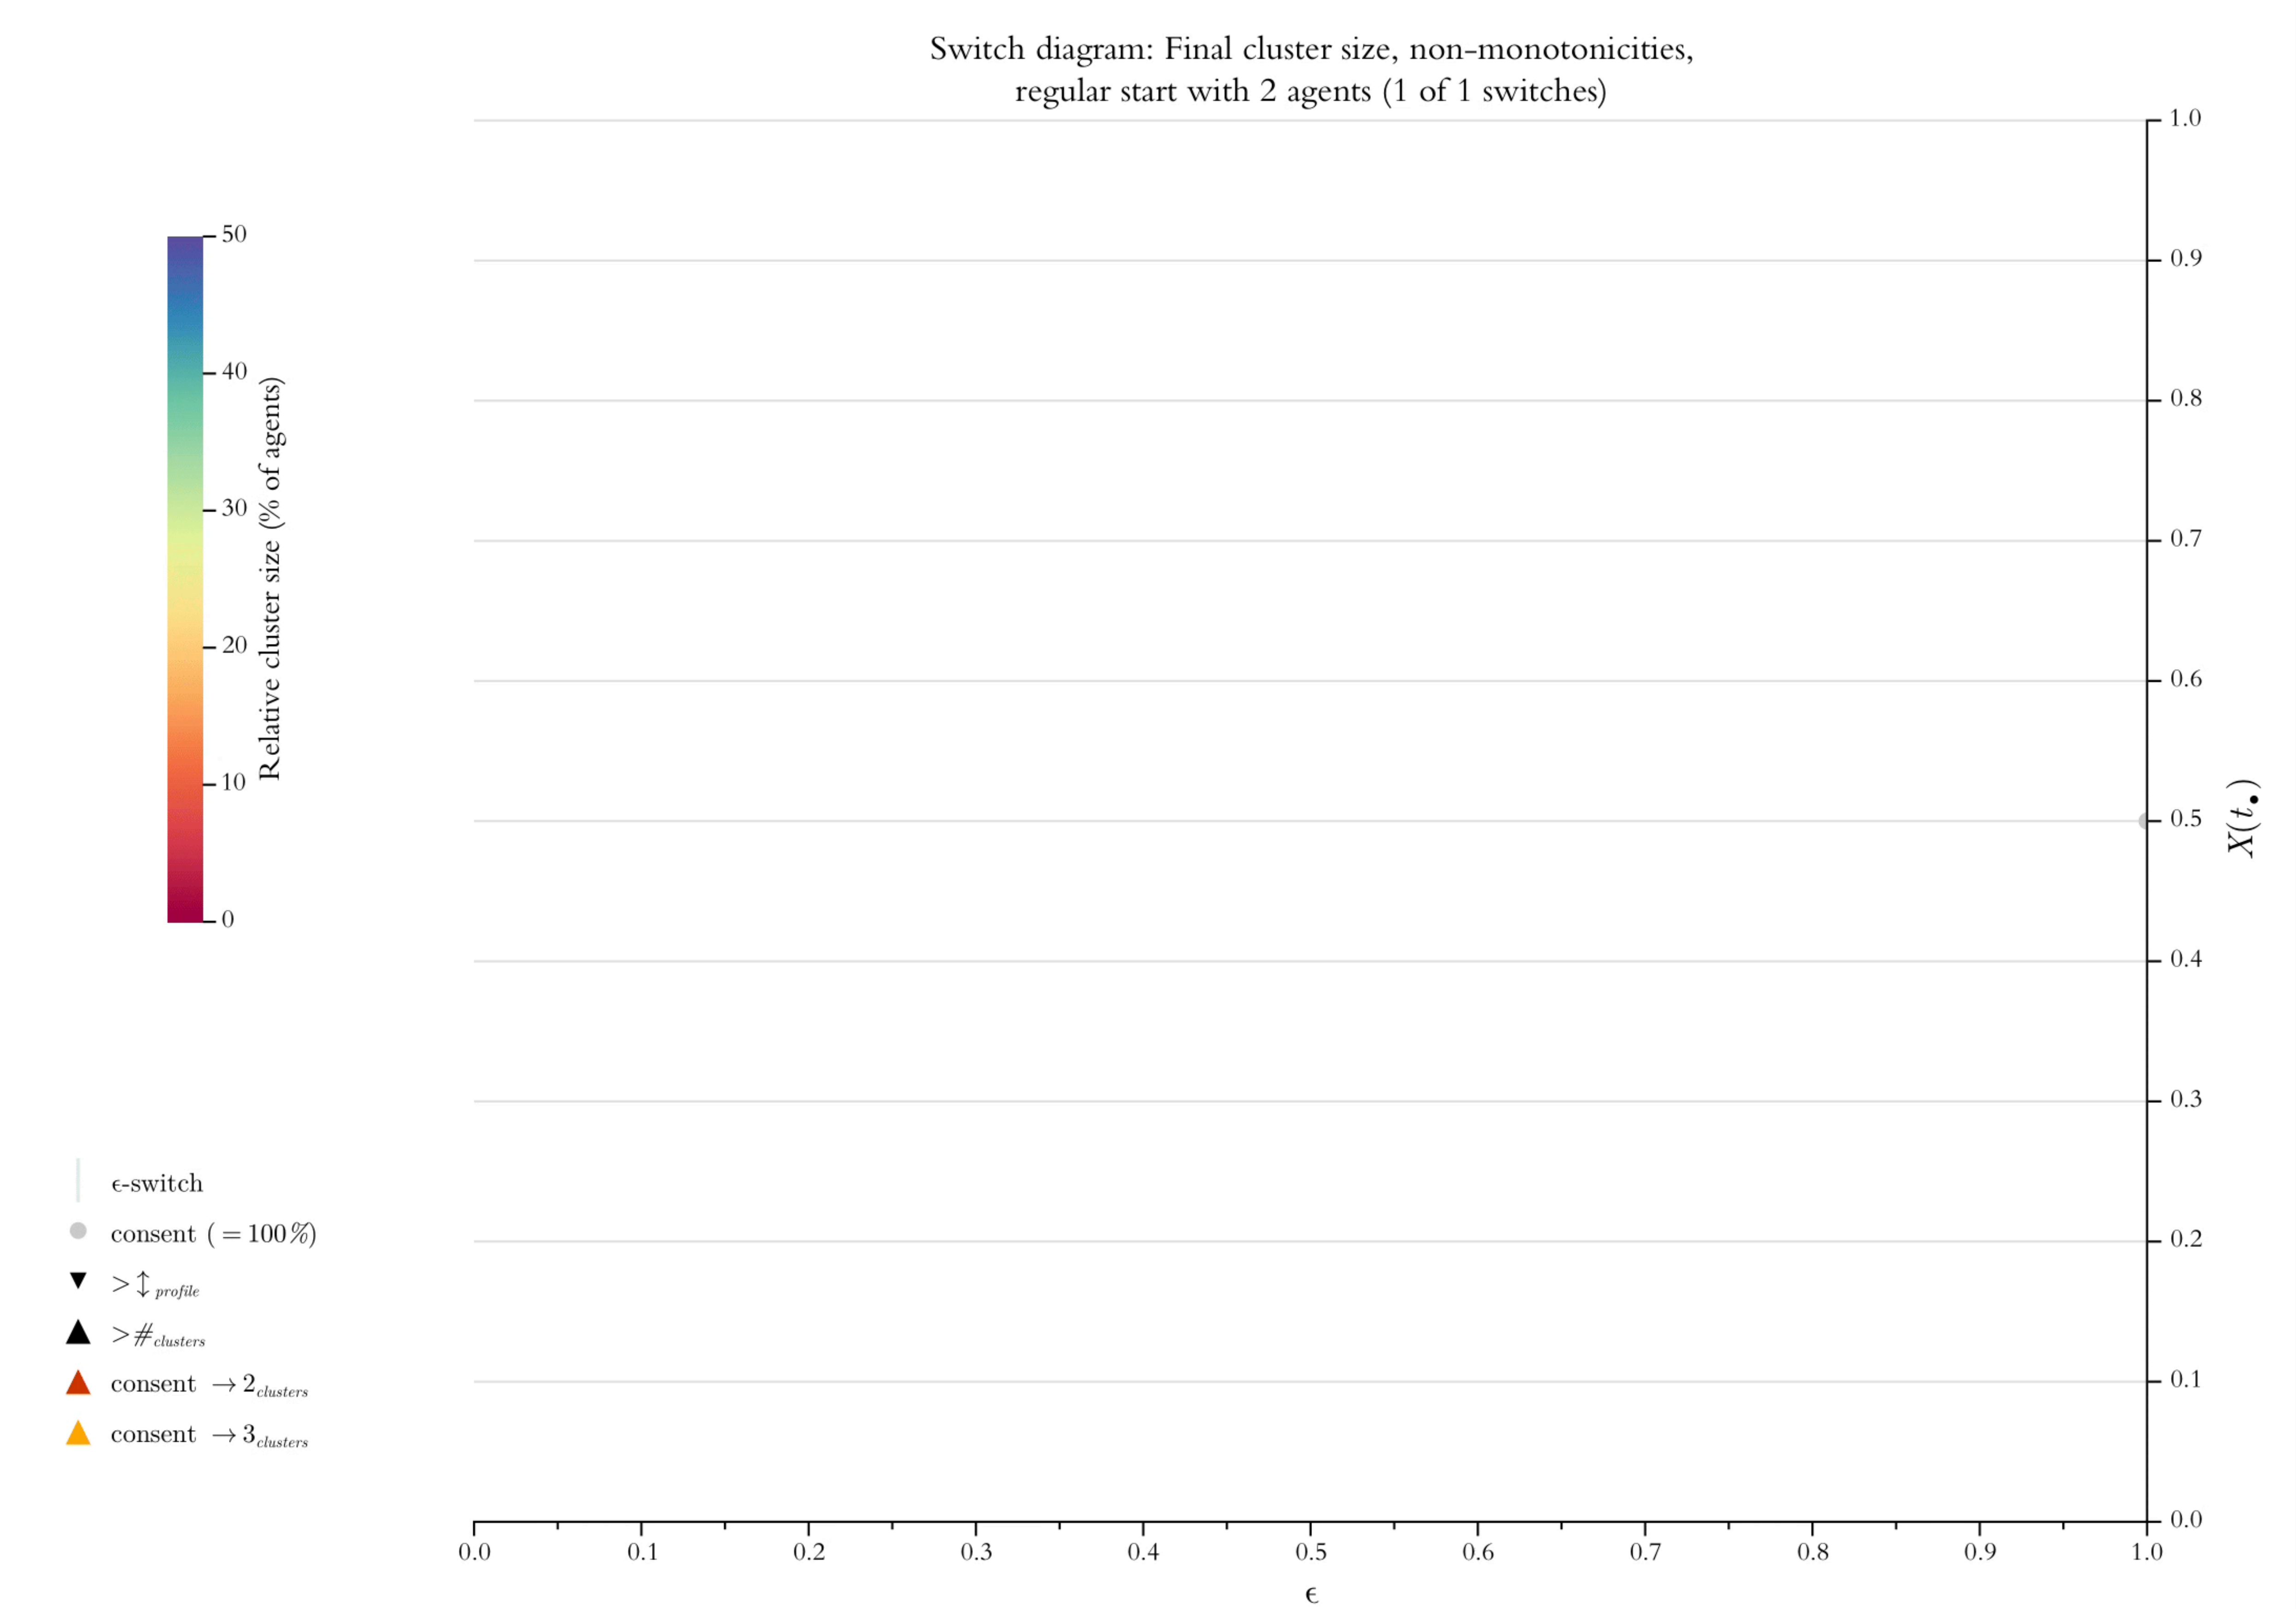

ϵ-*switch diagrams*: regular start, $n = 3,5,\ldots,79$

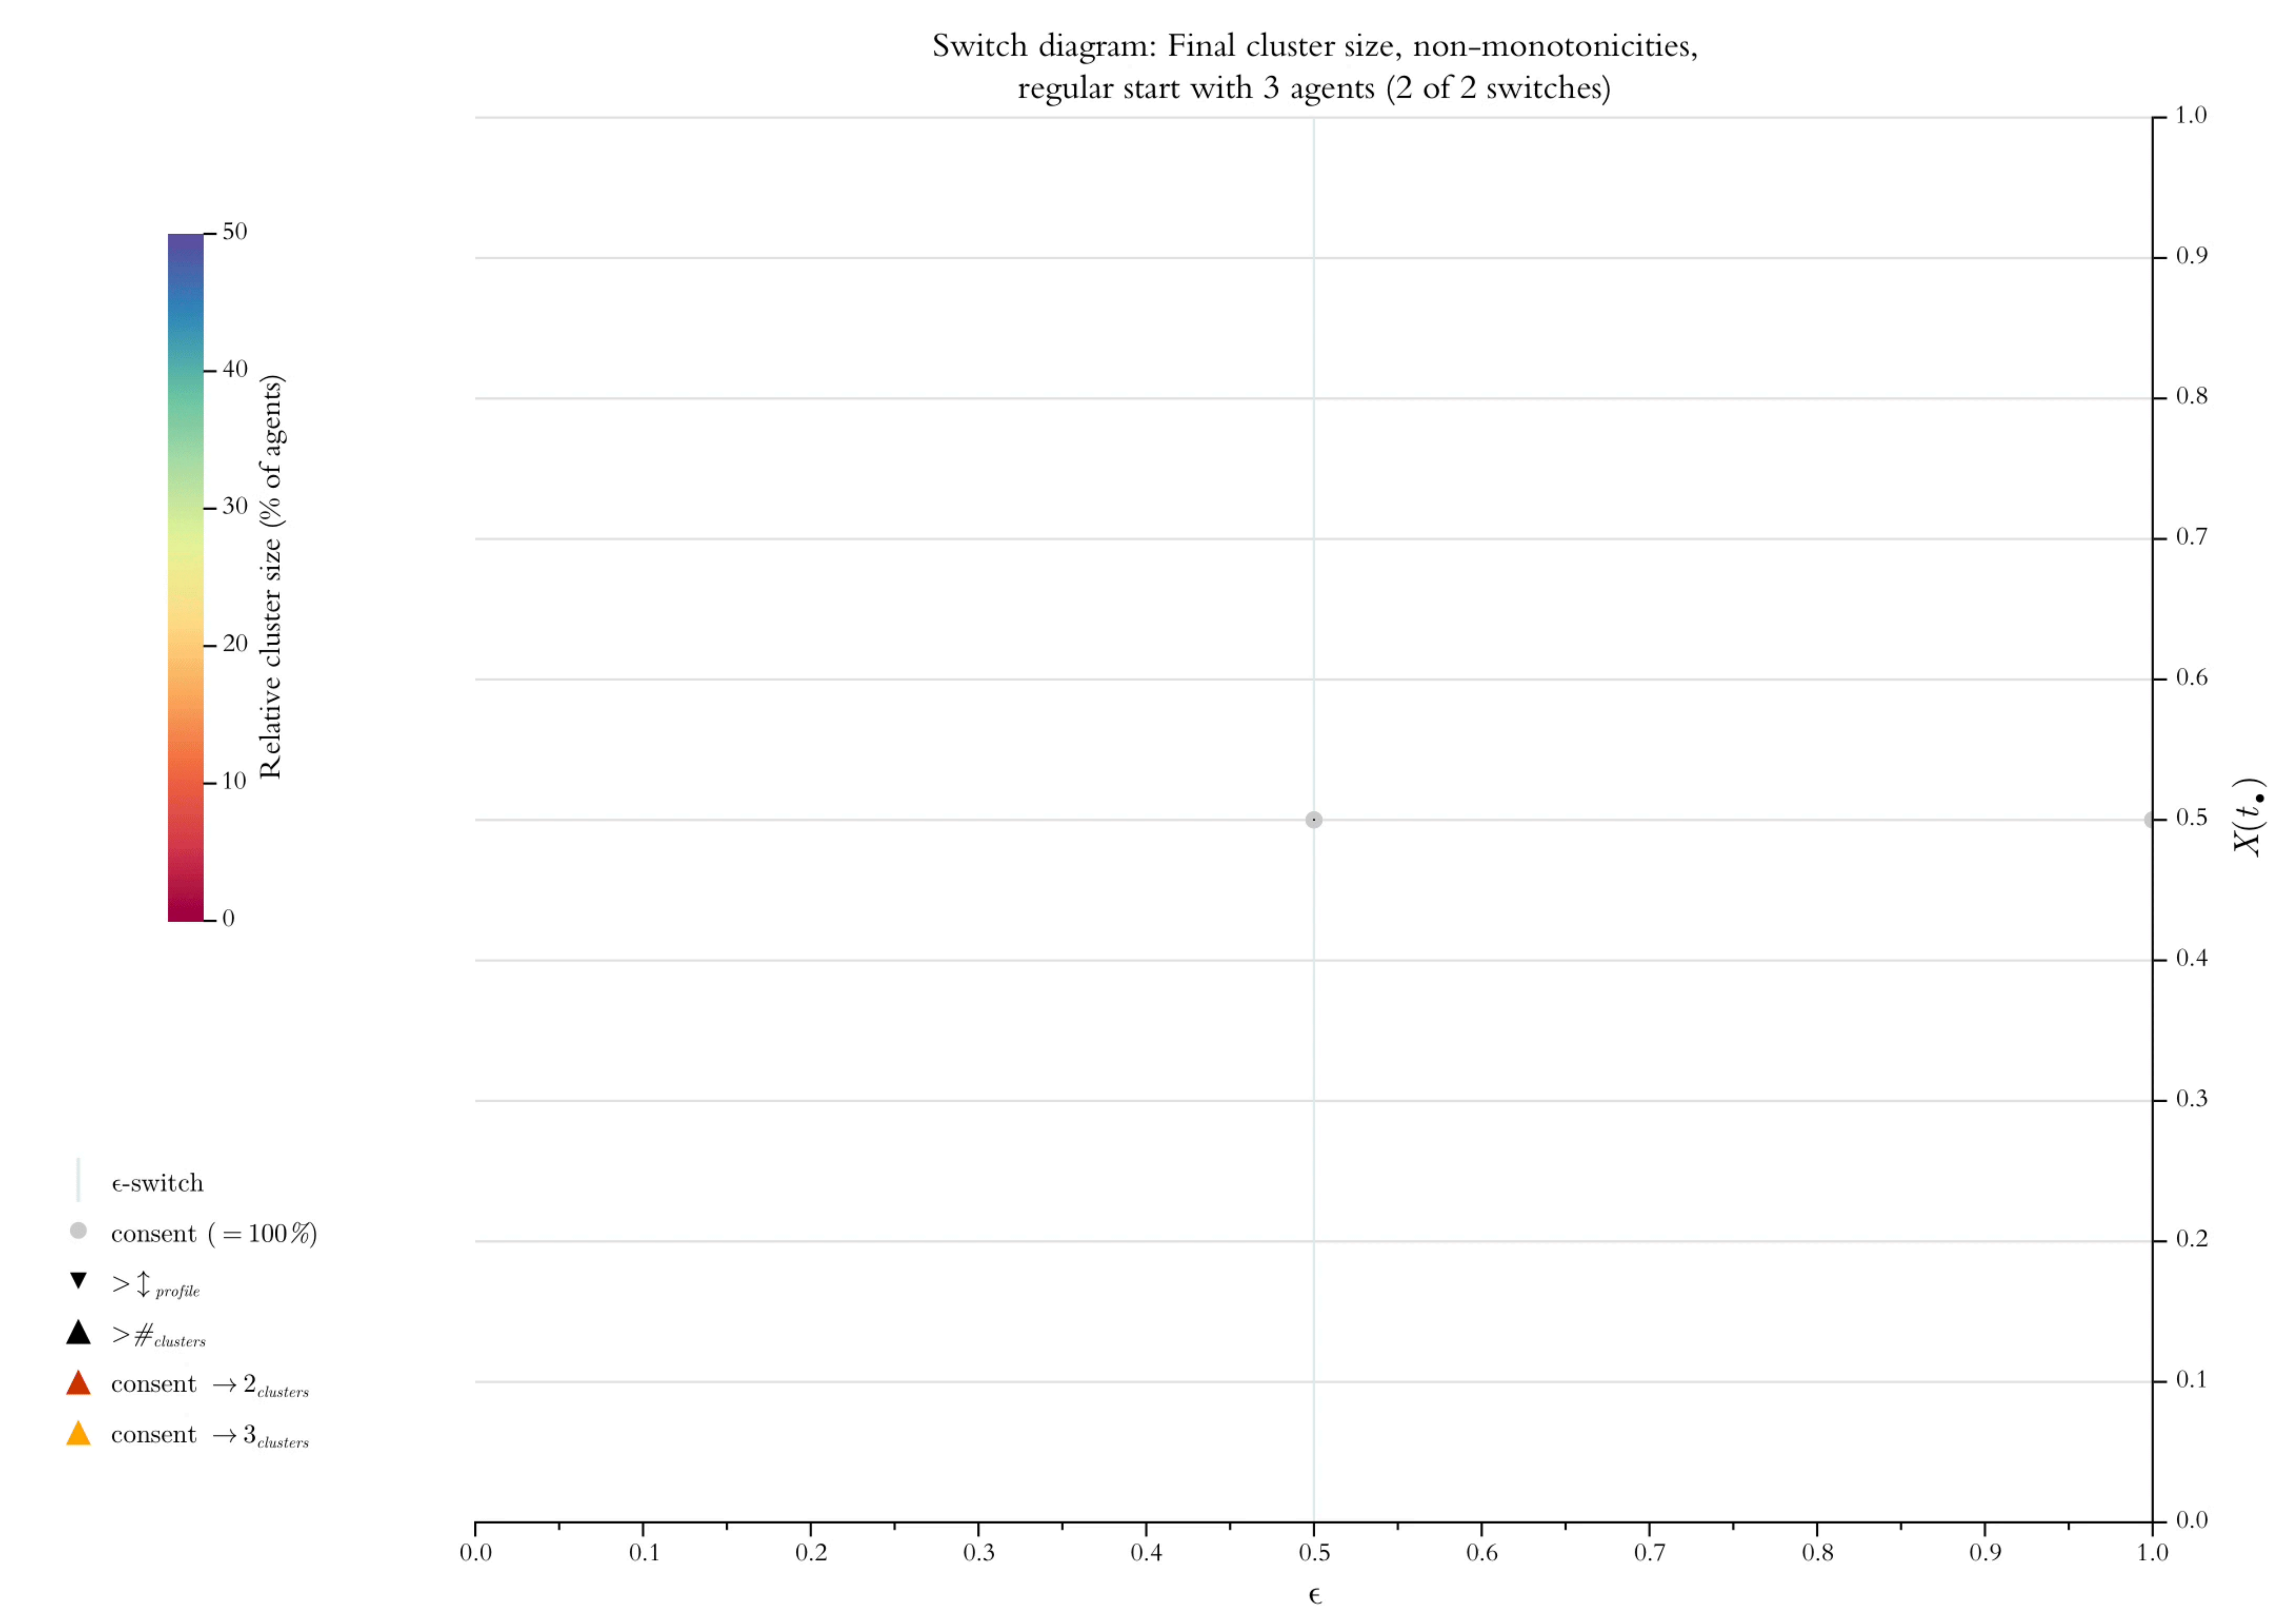

Sequences of ϵ-switch-based *single* runs

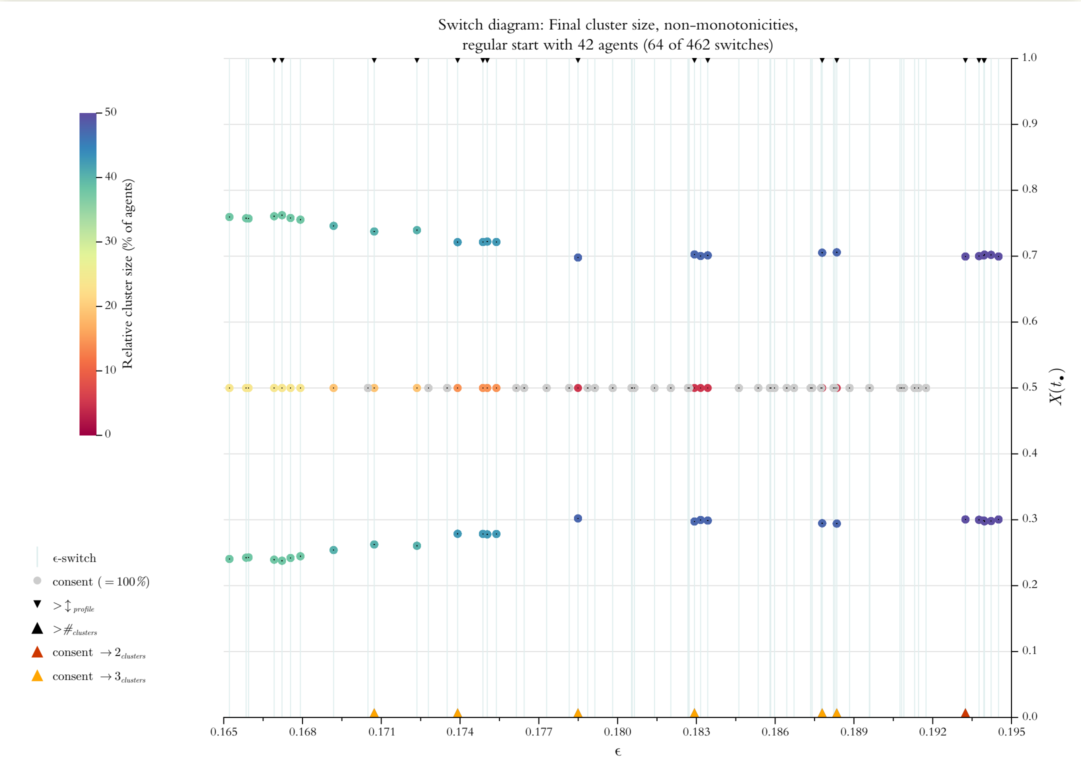

How Comes?

single runs

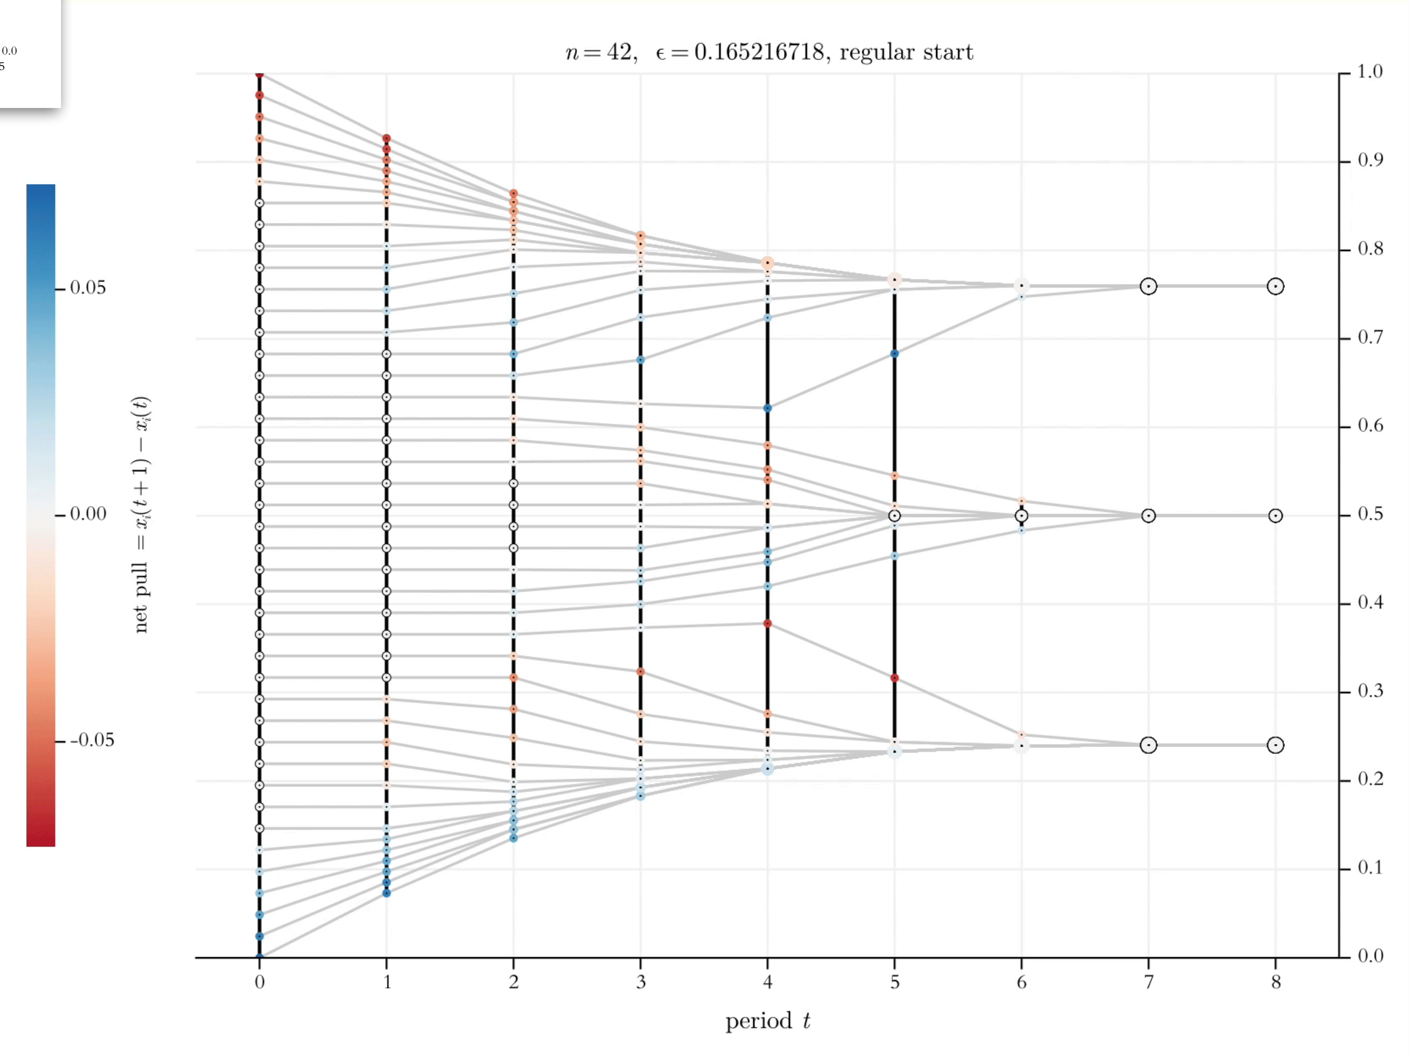

How Comes?
The BC-dynamics as a *dynamical* network

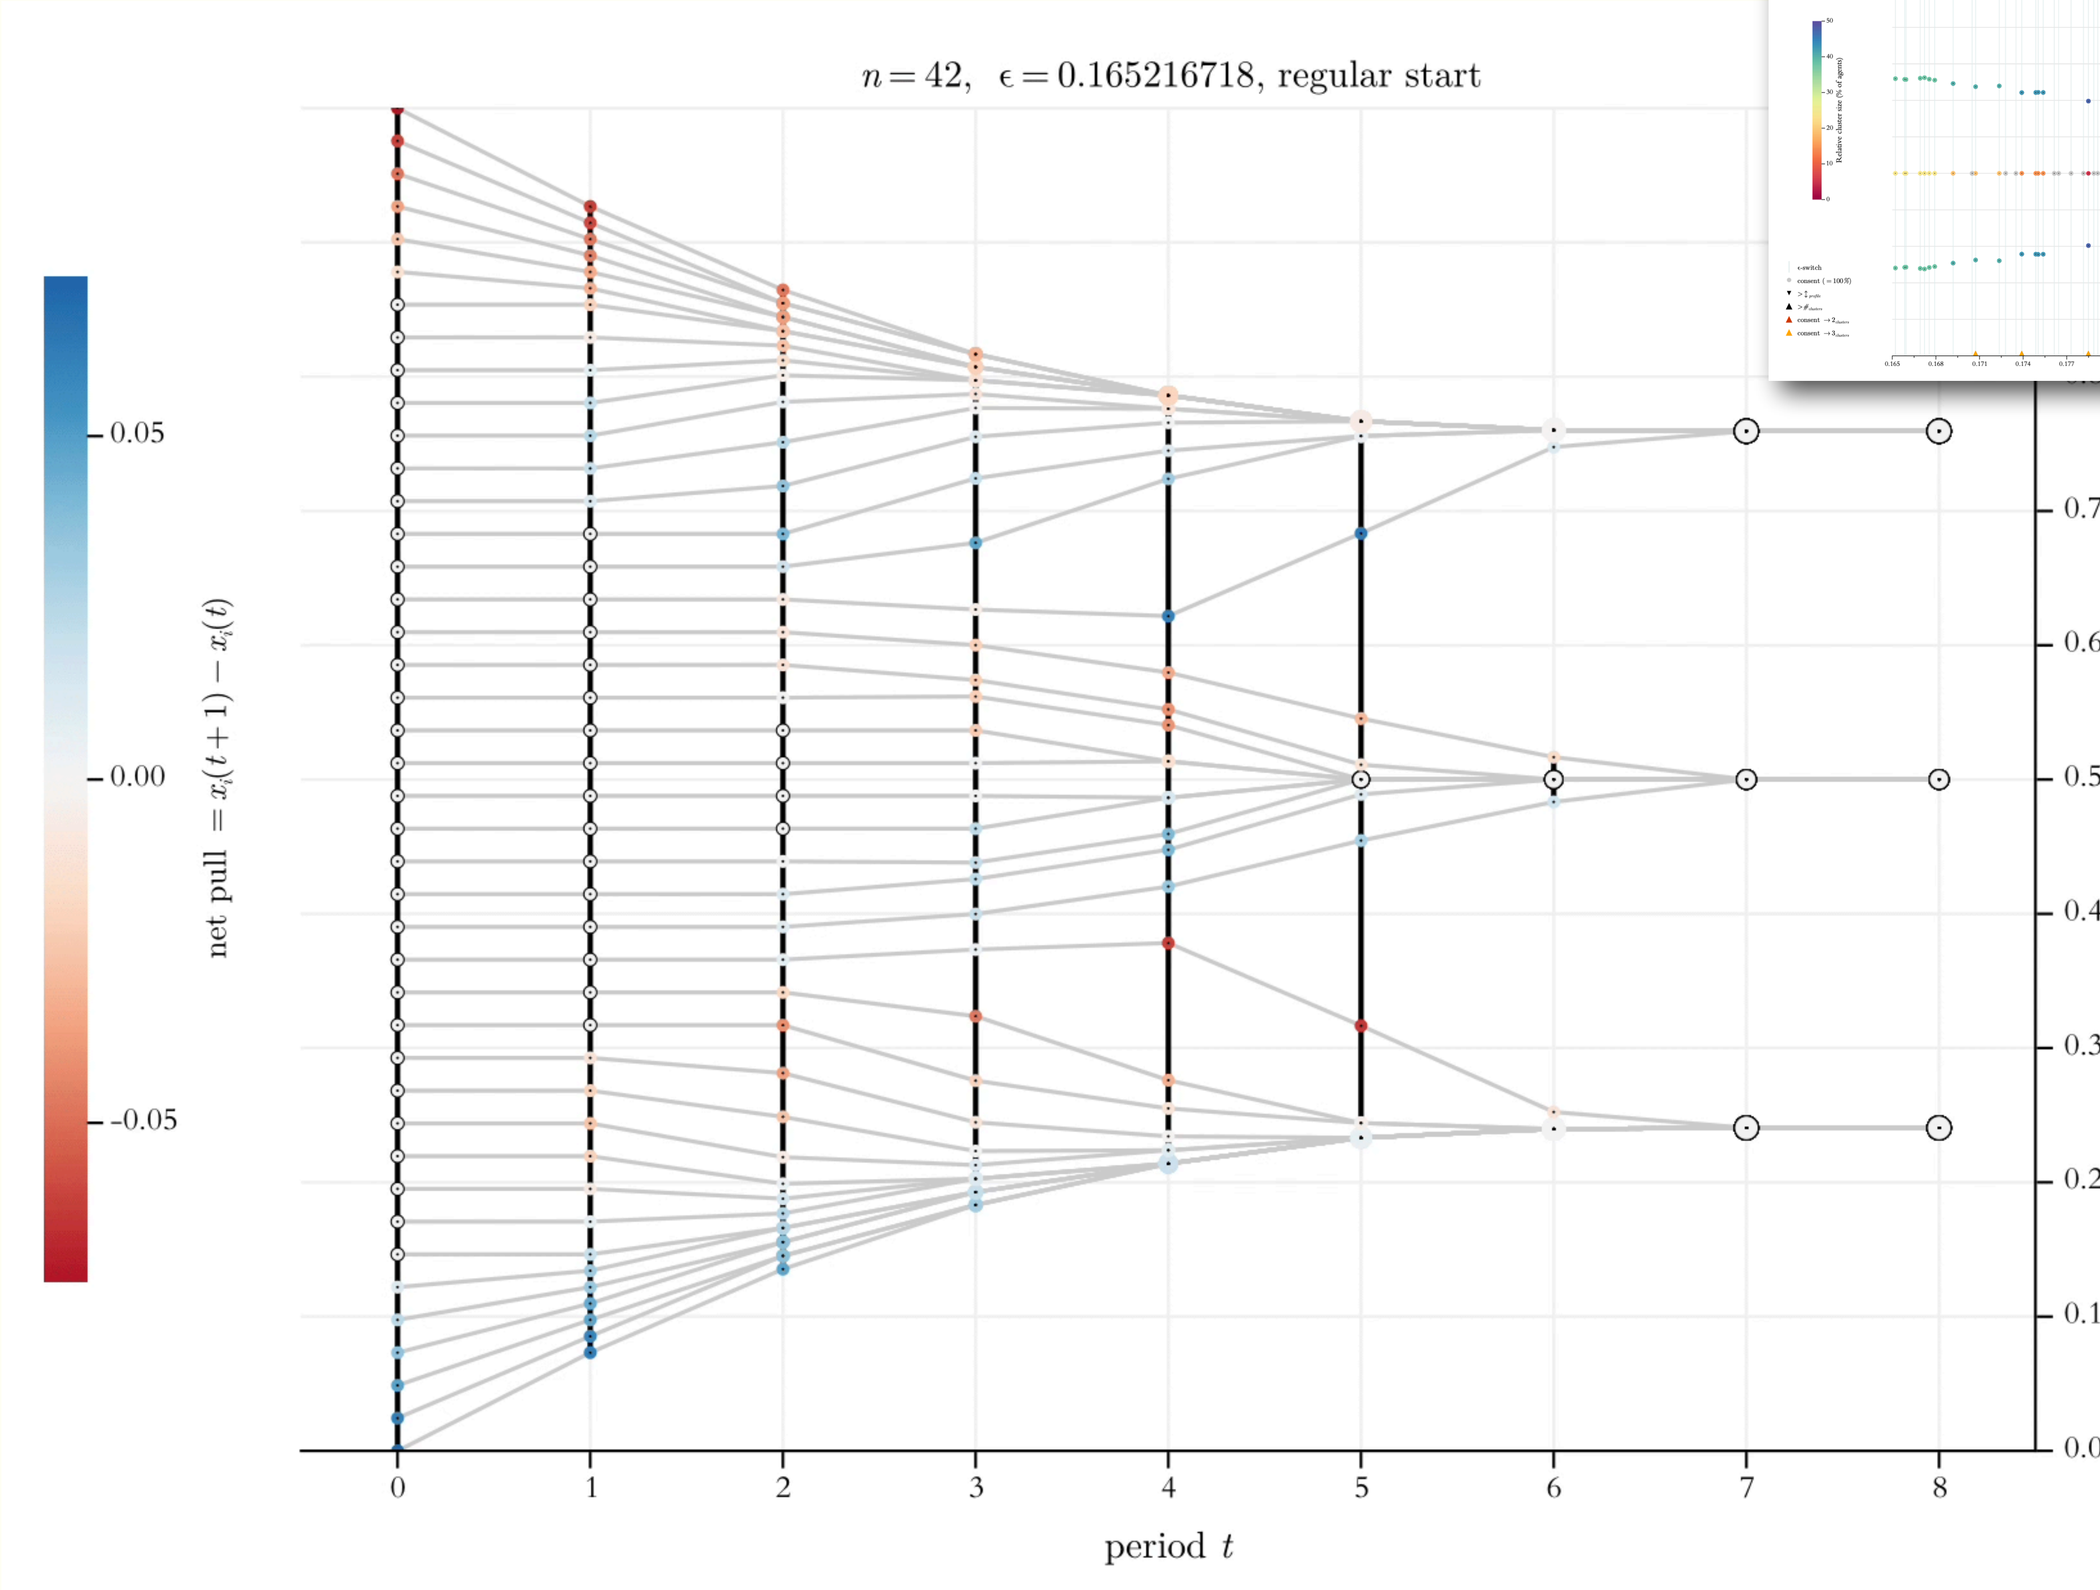

Switch diagram: equidistant start, $n = 42$

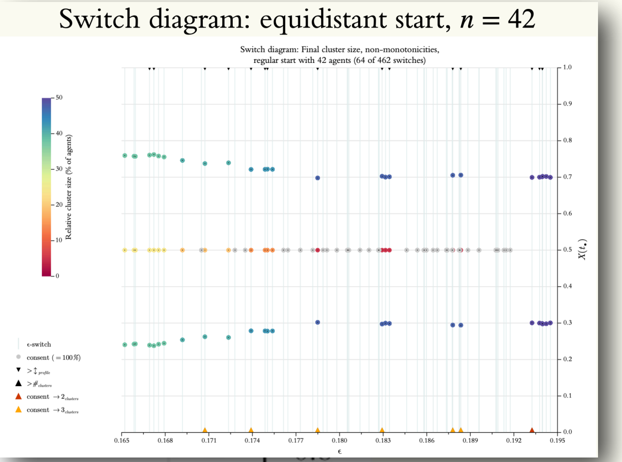

The key explanatory element:

Emergence and destruction of bridges in and between networks

First two of three *consecutive* switches

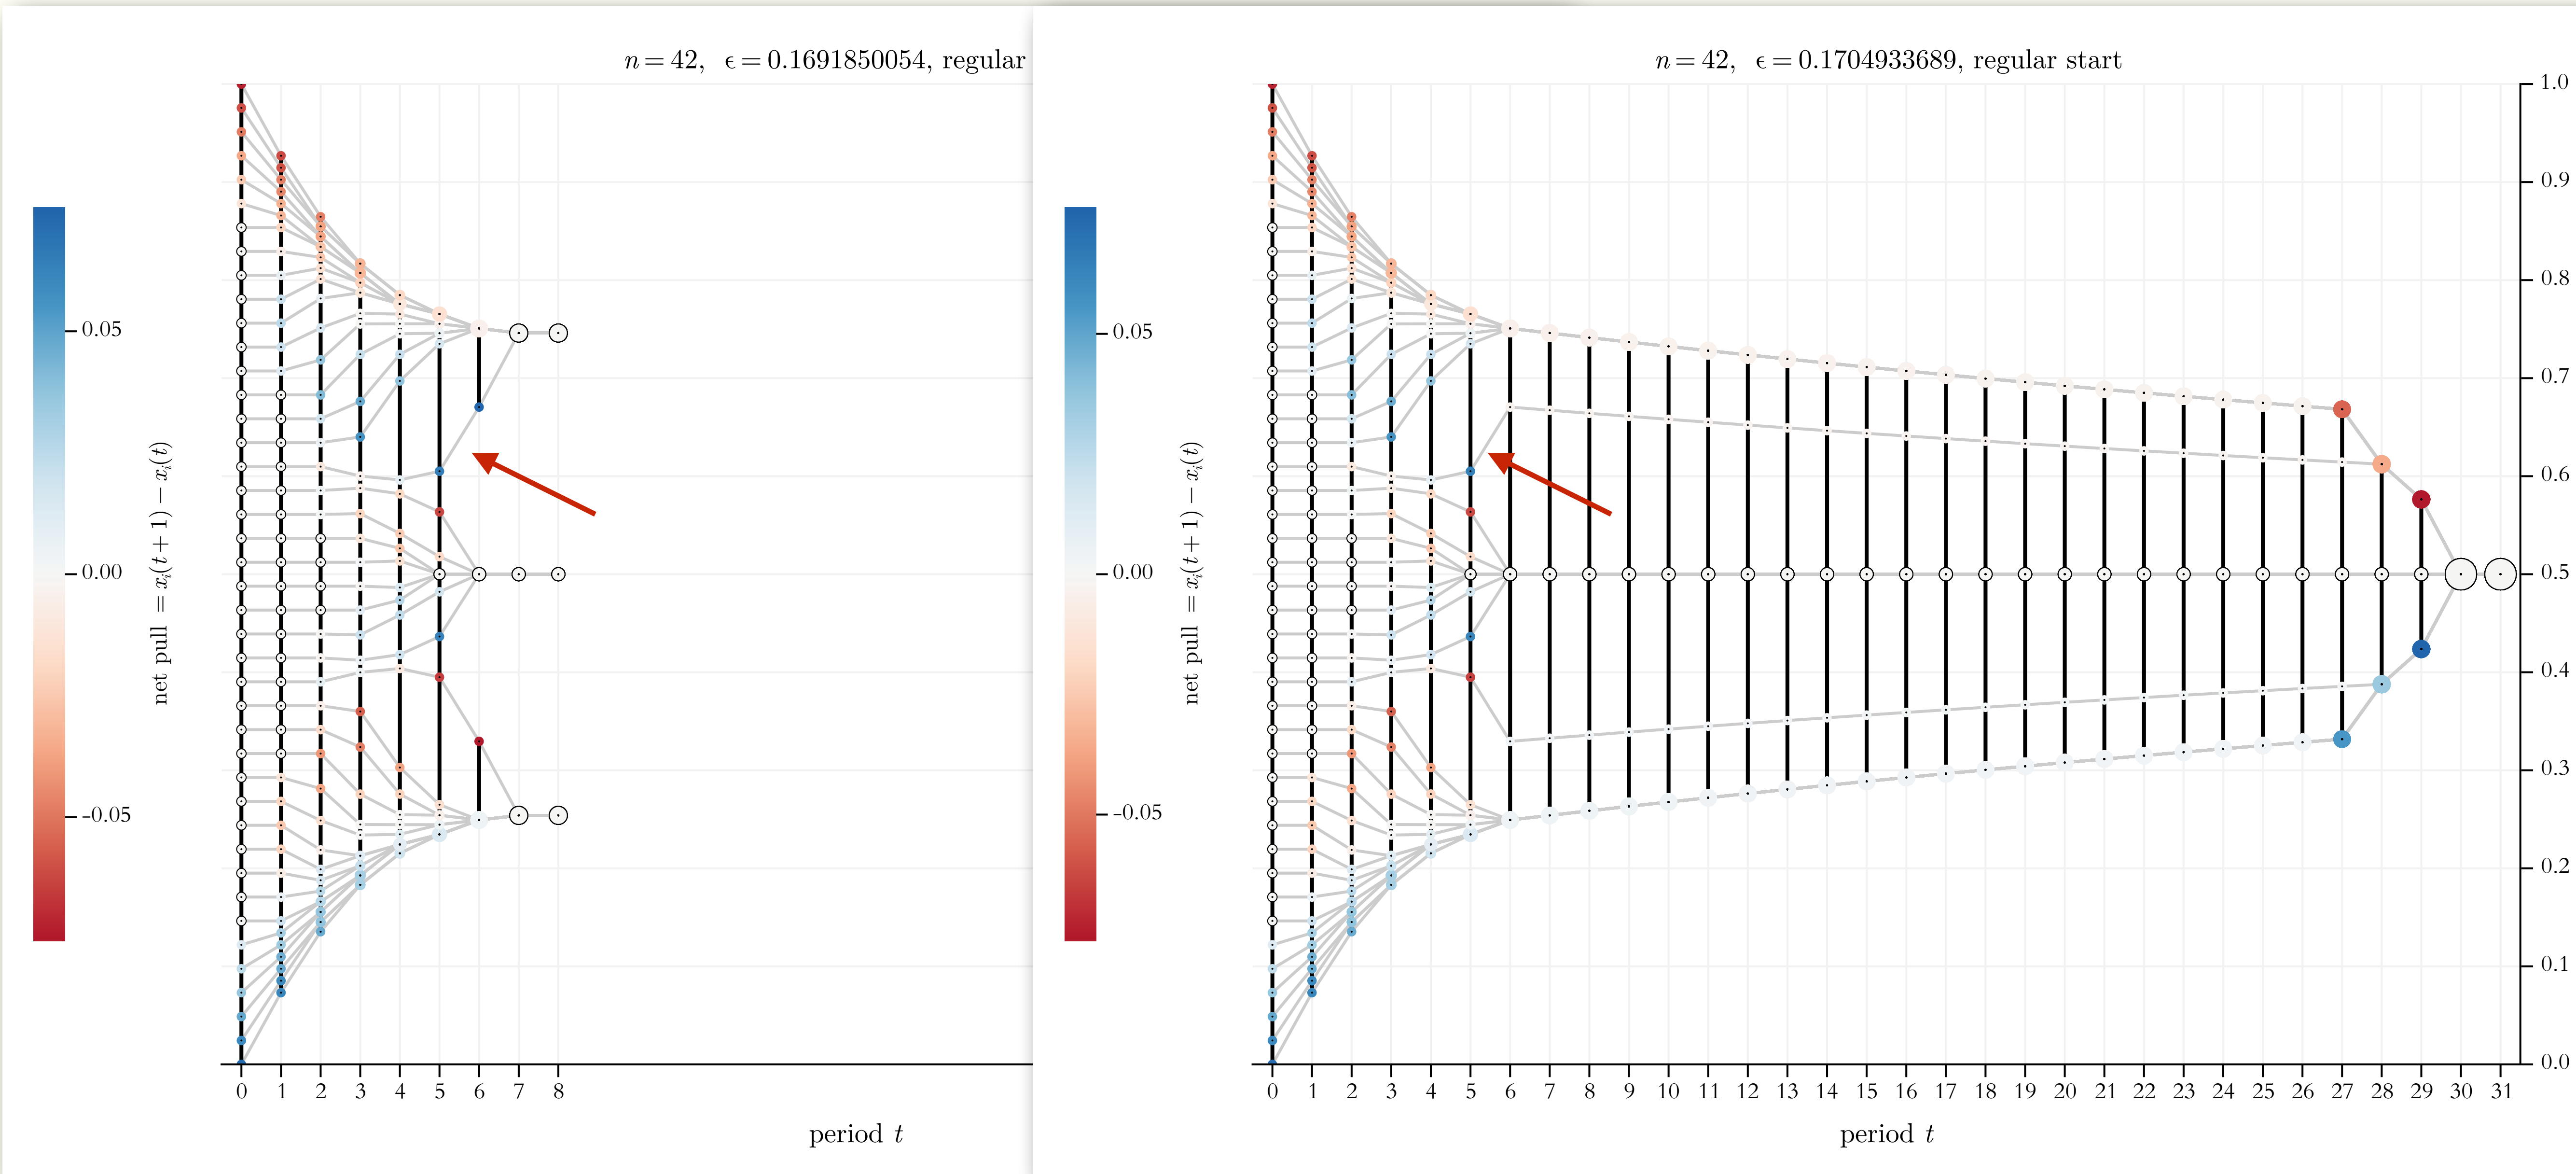

The key explanatory element:

Emergence and destruction of bridges in and between networks

Last two of three *consecutive* switches

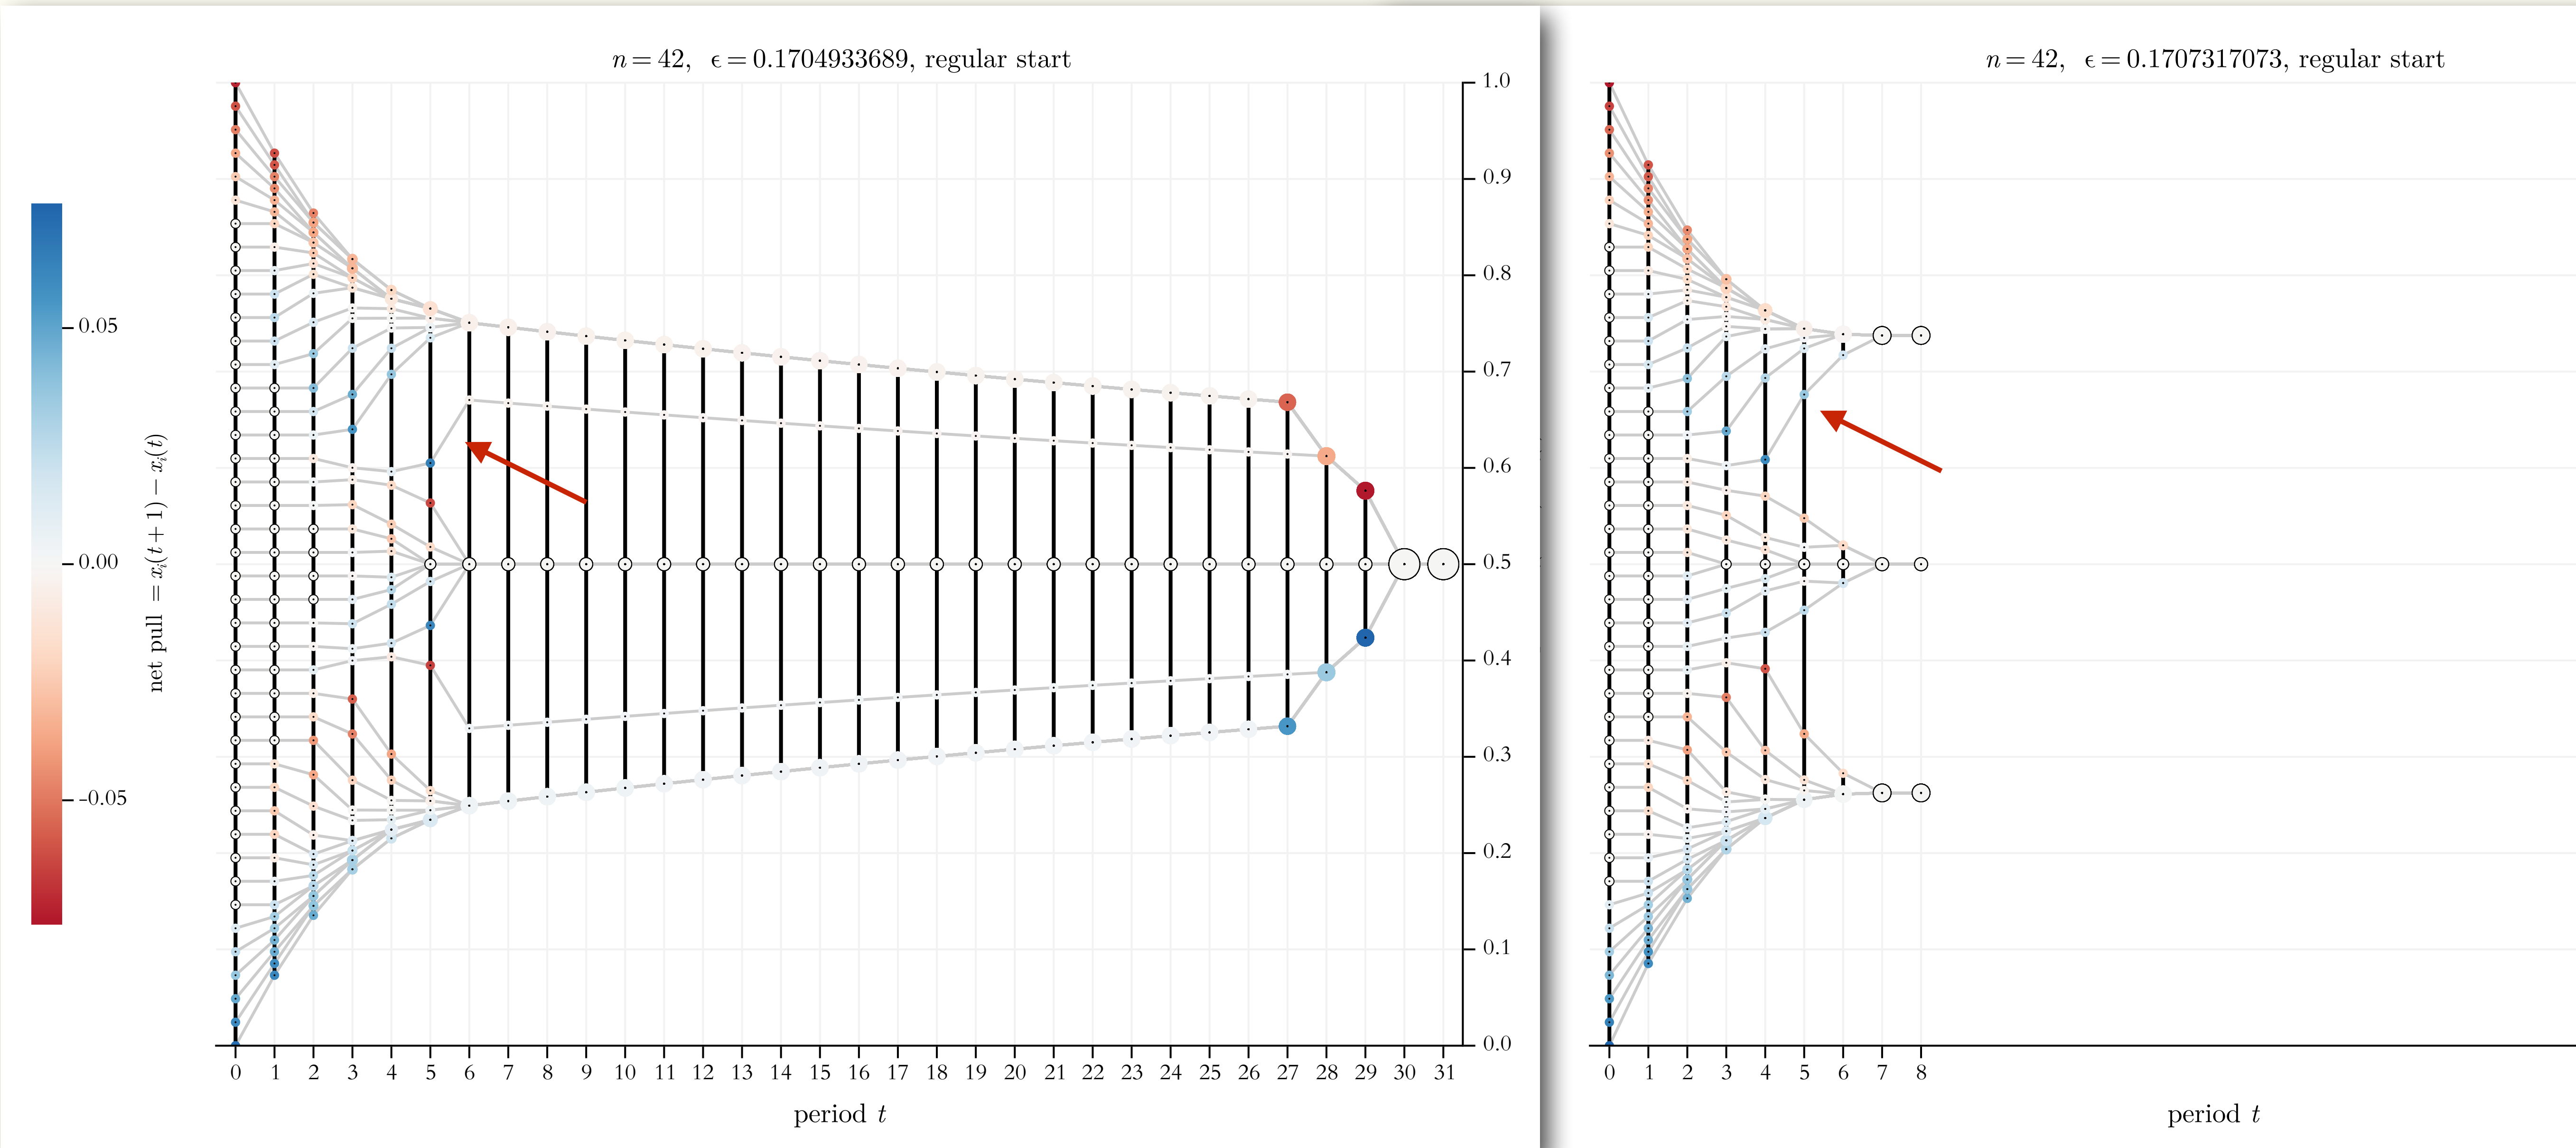

§5
Lessons and perspectives

Relevance & lessons

In general

- The effects are *not due to the equidistance.* [In 100 experiments with uniform random starts (always 40 agents) I found 65 cases with two-cluster polarisation after a consensuses, in 37 cases multiple times. However, I have found no case of three cluster-polarisation after a consensus.]
- Noise or heterogeneity may somehow and to some degree *smooth out* the effects. But to fully understand what noise or heterogeneity does, *one must first understand the idealised (contra-factual) processes*.

If a real world opinion dynamic is predominantly a BC-process, **then**

- whether or not we get polarisation may be *highly sensitive to initial conditions.*
- „Be more open-minded!" is *not an advise that works against polarisation under all conditions*; it might be a recipe to produce polarisation.
- *policy recommendations are extremely difficult,* success & failure may be a matter of luck.
- a *new type of study* is necessary: *Robustness of policies* in the face of *imprecise knowledge* about a society's position in a parameter space.

Living with a model – some final confessions

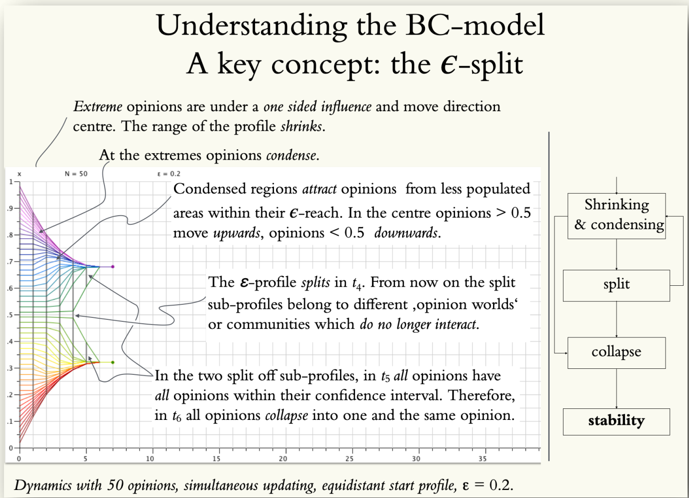

stage 1

I found the model a bit boring, too easy to understand, too easy to reflect on its status, now understood, no future perspective, in short: *nothing to fall in love with.*

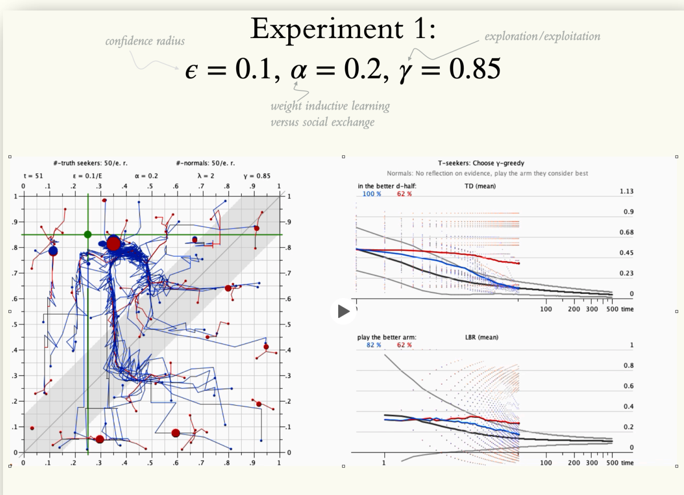

stage 2

I realised that it was just the simplicity of the model that made it an ideal starting point to do more interesting things with it.

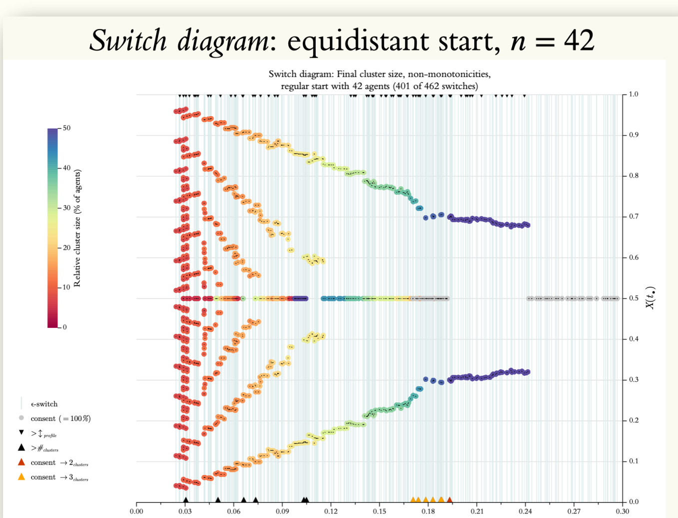

stage 3

I realised, I had underestimated the model for more than 15 years.

There is much more complexity in it than I had thought at the beginning.

A surface simplicity conceals a beauty (of complexity).

I have fallen in love with the model—finally:-).

Many thanks for your attention!