Random Walk Processes in ACT-R Mechanisms Lead to a Wild Distribution of Learning Times

Richard M Young
Anna Cox
and Martin Greaves

Psychology Department
University of Hertfordshire

This research was supported in part by the US Air Force, European Office of Aerospace Research & Development, Contract no. F61775-01-W-E040.

Variability in Concept Learning Times

- We have noticed, in simple Act-R models of concept learning, a high degree of variability in learning times.
- This table shows the standard deviation between runs of "experiment(N)" for learning in a simple model, i.e. the SD between estimates averaged over N simulated subjects.

N	Mean # trials	SD of mean
5	18.19	4.01
10	17.06	2.95
20	16.82	1.98
50	17.18	1.40
100	16.95	0.88
200	17.01	0.62
500	17.16	0.41

- Those large SDs, e.g. for experiment(50) or experiment(100), make it difficult to estimate stable fits to data,
  - — e.g. for optimising parameters, or plotting the "landscape" in parameter space.

The "Fast" Kendler & Kendler Model

- Our examples are taken from a version of Niels Taatgen's model of the "fast" Kendler & Kendler (1959) children, as described in Anderson & Lebiere (1998) book.

- The task involves children being presented with objects which differ on two binary dimensions:

  colour: {red, green}
  size: {big, small}

- Initially, the concept is "red" and Ss are supposed to say YES to red, NO to green. Once they have it right (criterion: 10 consecutive correct), in one condition the concept is secretly switched to "big", and they have to say YES to big and NO to small — ignoring the colour.

- Key part of the model for us is that it contains two productions, attend-color and attend-size, which compete and of which exactly one fires.
  - If the correct production fires, it leads (after rapid learning) to 100% success.
  - If the wrong production fires, it leads to ~50% success.

Long Learning Times

- The high SD of the mean learning times, as reflected in the "experiment(N)" results,
  - — implies that there is a high SD for the individual learning times;
  - — but also suggests the possibility that there is a "wild" distribution of learning times, i.e. with some individual very long learning times. Although these would be rare, even one of them in a sample would be enough to disturb the mean learning time.
- If we look at the distribution of learning times, the overall shape is reasonable, but the curve flattens out "too much", i.e. falls too slowly at high values.
  - — we are currently engaged in some curve-fitting to try to make this claim more precise.
- Where might this wild distribution come from?

Distribution of Learning Times for “Non-Reversal” Condition

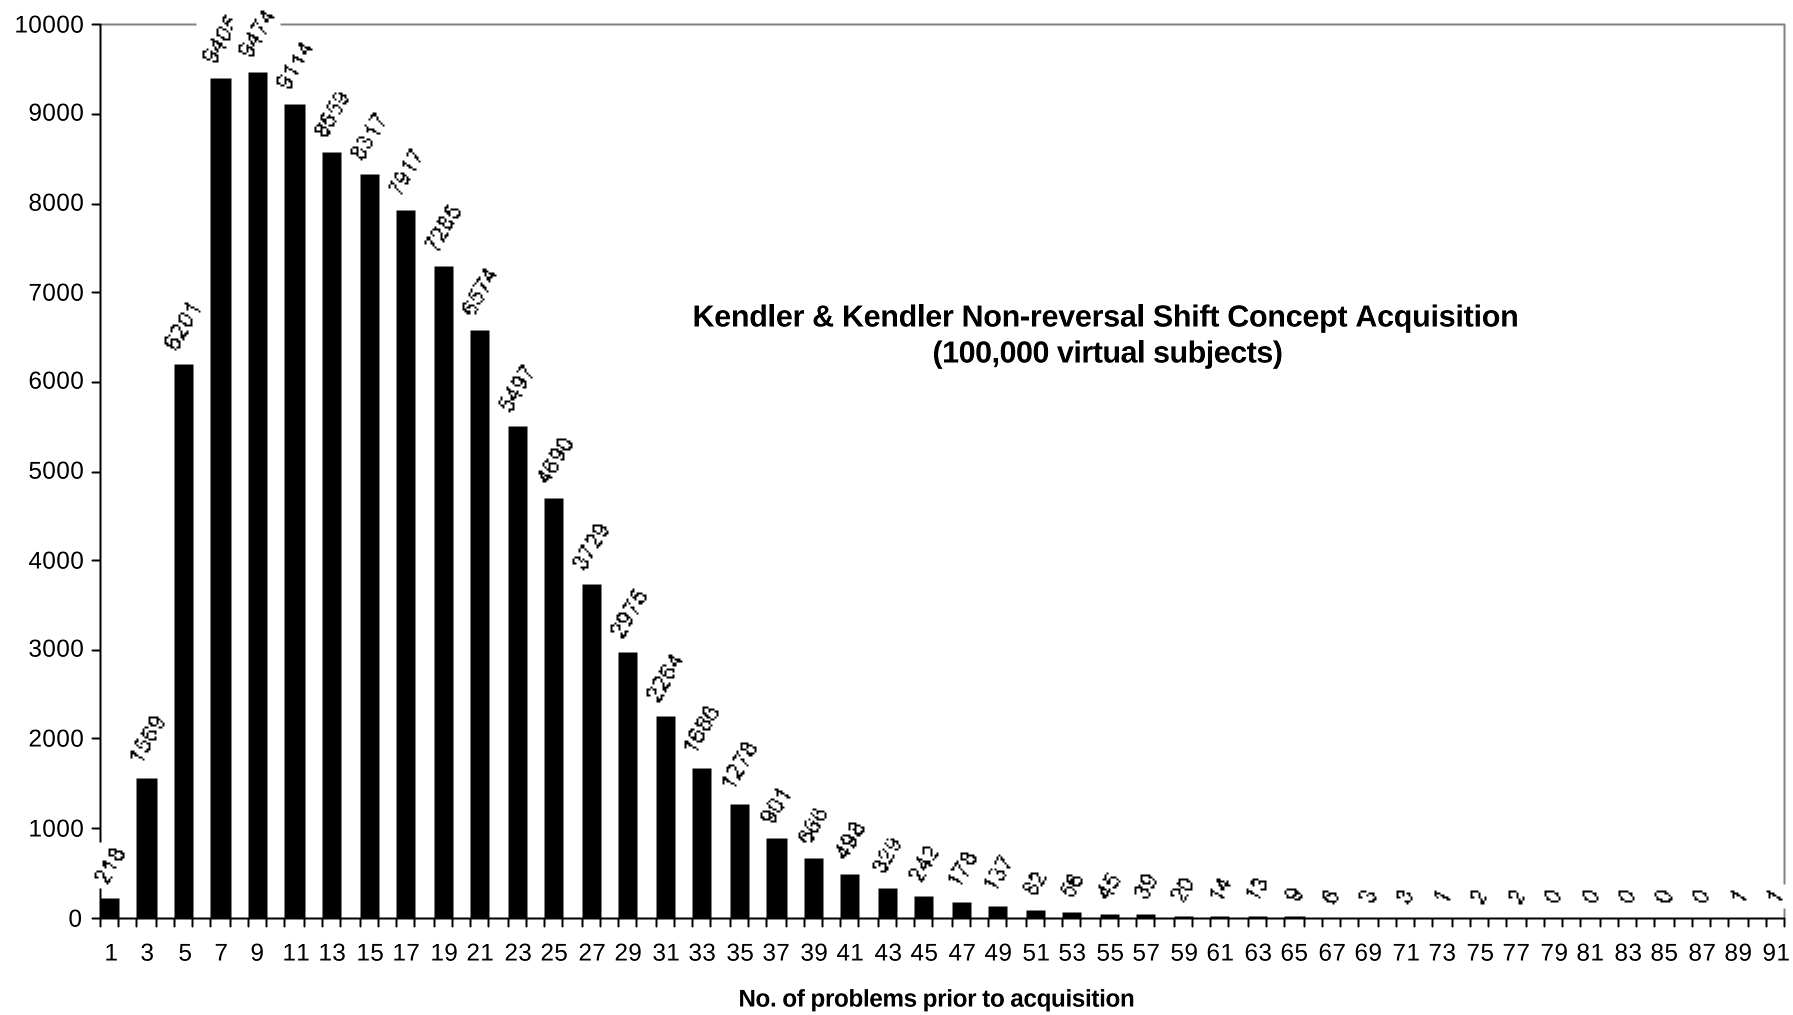

Random Walks

- Another process which generates wild distributions is a random walk. In the simplest case, the symmetric random walk, have a variable which randomly adds ±1 at each step.

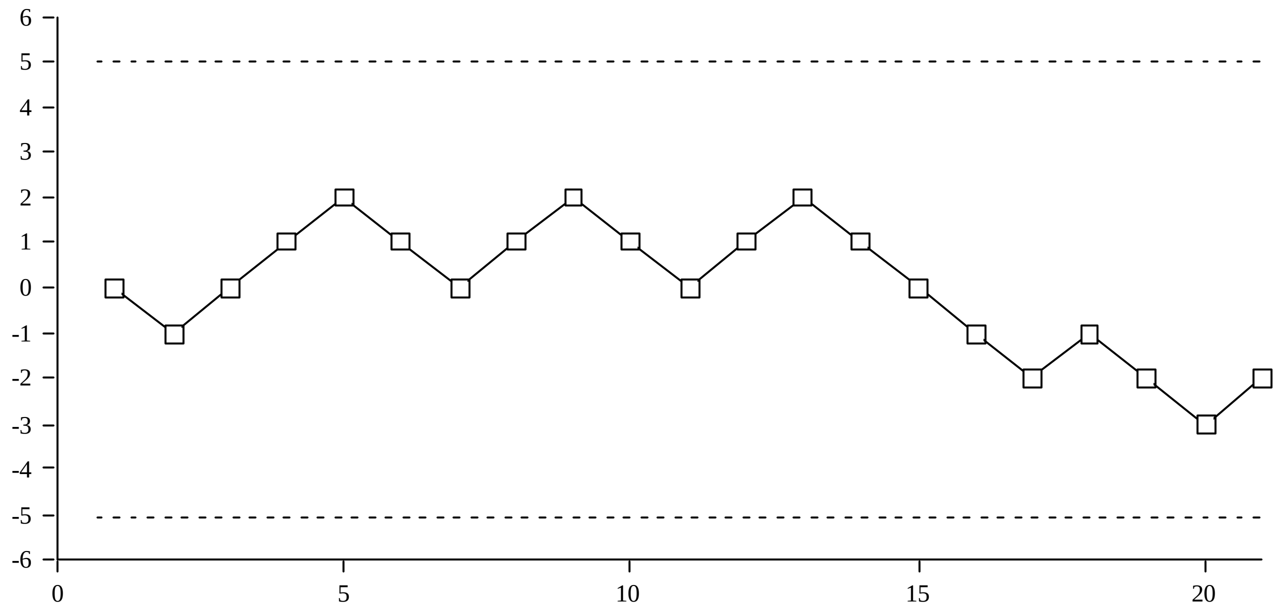

- Interesting properties of random walks concern their *times to first passage*. The picture has barriers at ±5. The mean time to reach one or other of those barriers is $5^2 = 25$.

- But the time to reach a specified one of the barriers, say +5, has a wild distribution — so wild that its mean is infinite!
  - — The modal time to reach the +5 barrier is 7-9 steps, but the median time is 53 steps, the third quartile is not reached until 245 steps, while the 90-percentile is greater than 1200. In other words, 10% of the walks take more than 1200 steps to reach the +5 barrier.

Distribution of Random Walk Steps to First Passage (single barrier)

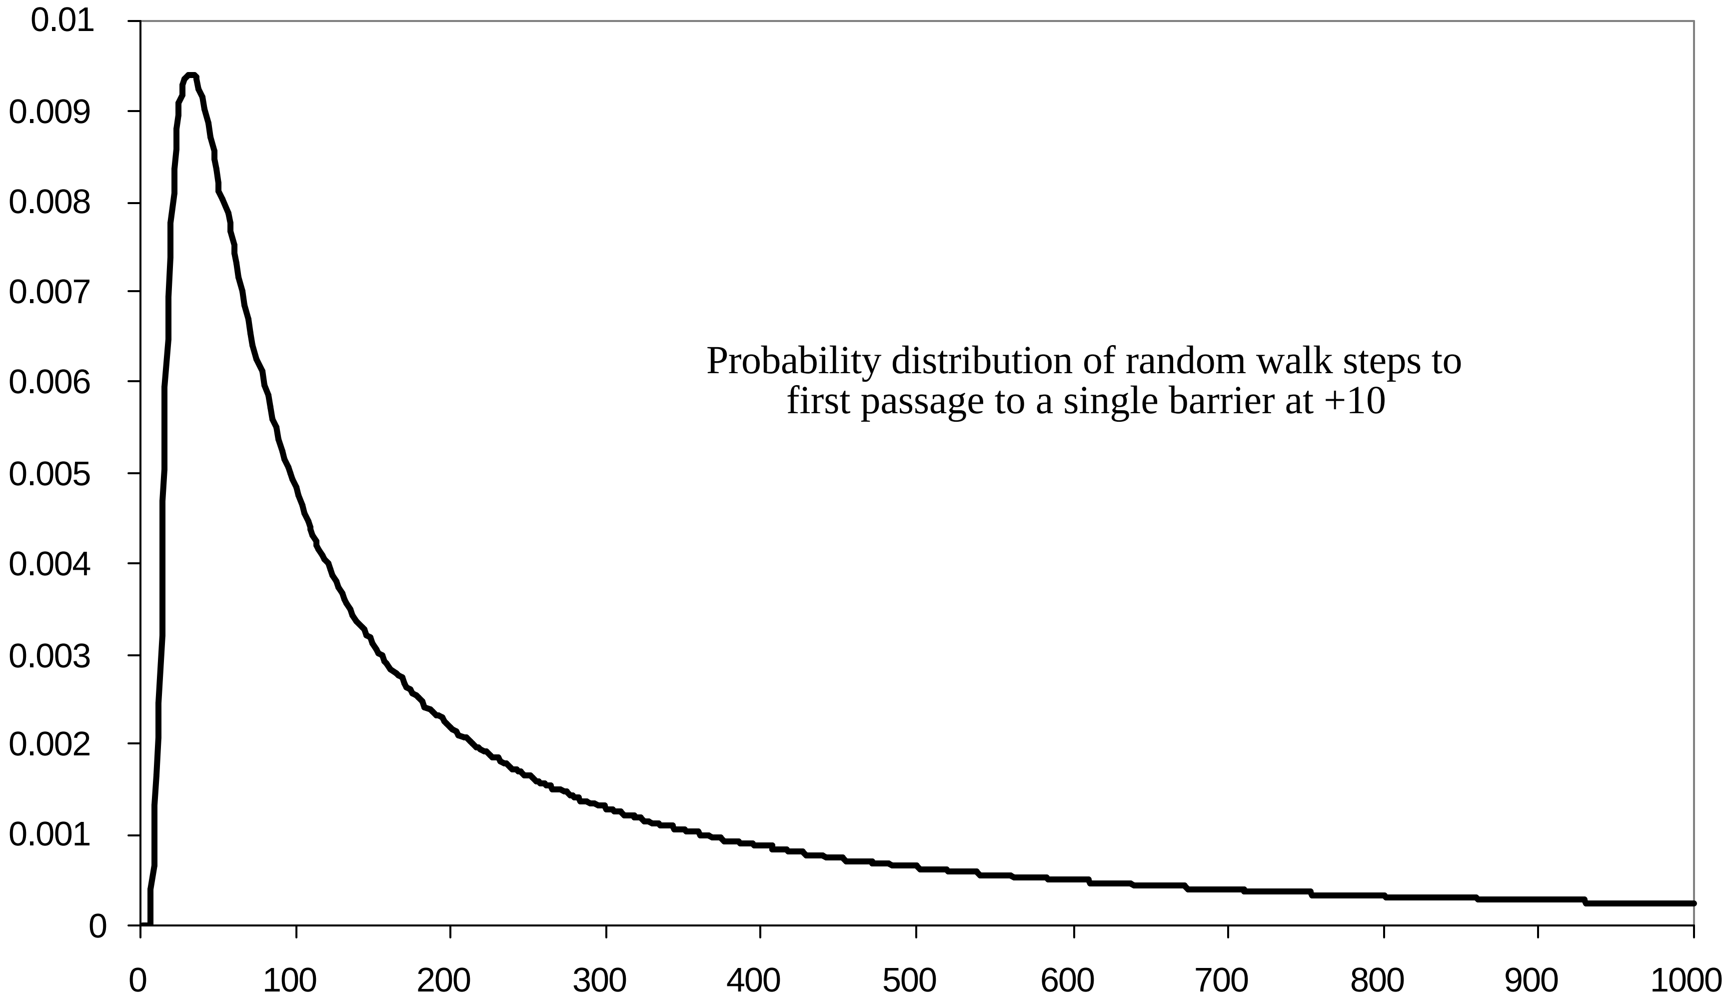

Random Walks in Act-R Learning

- Do random walks arise in the normal Act-R learning mechanisms? Yes, they can do.

- Consider the *attend-color* and *attend-size* productions. Remember that the chance of each production being chosen depends on its PG–C, which in turn depends on its *r*, reflecting the proportion of times it has led to success.

- Immediately following the initial learning, *attend-color* will have a strong history of success, including the 10 successes in a row. Its PG–C is (usually) sufficiently more than *attend-size*'s that it fires all the time: *attend-size* doesn't get chosen.

- After the shift, *attend-color* experiences only ~50% success. Its PG–C therefore falls towards 0.5, but long before it gets there, its PG–C becomes comparable with that of *attend-size*. *Attend-size* therefore begins to fire, rapidly comes to gain 100% success, and soon takes over entirely from *attend-color*.

- But while *attend-color* is still in control, because half the time it still leads to success, the progress of its PG–C downwards follows a random walk.

  — so, some very long times before *attend-size* takes over

— also, luck (good or bad) can play a role.

Difference of PG–C between *attend-size* and *attend-color*

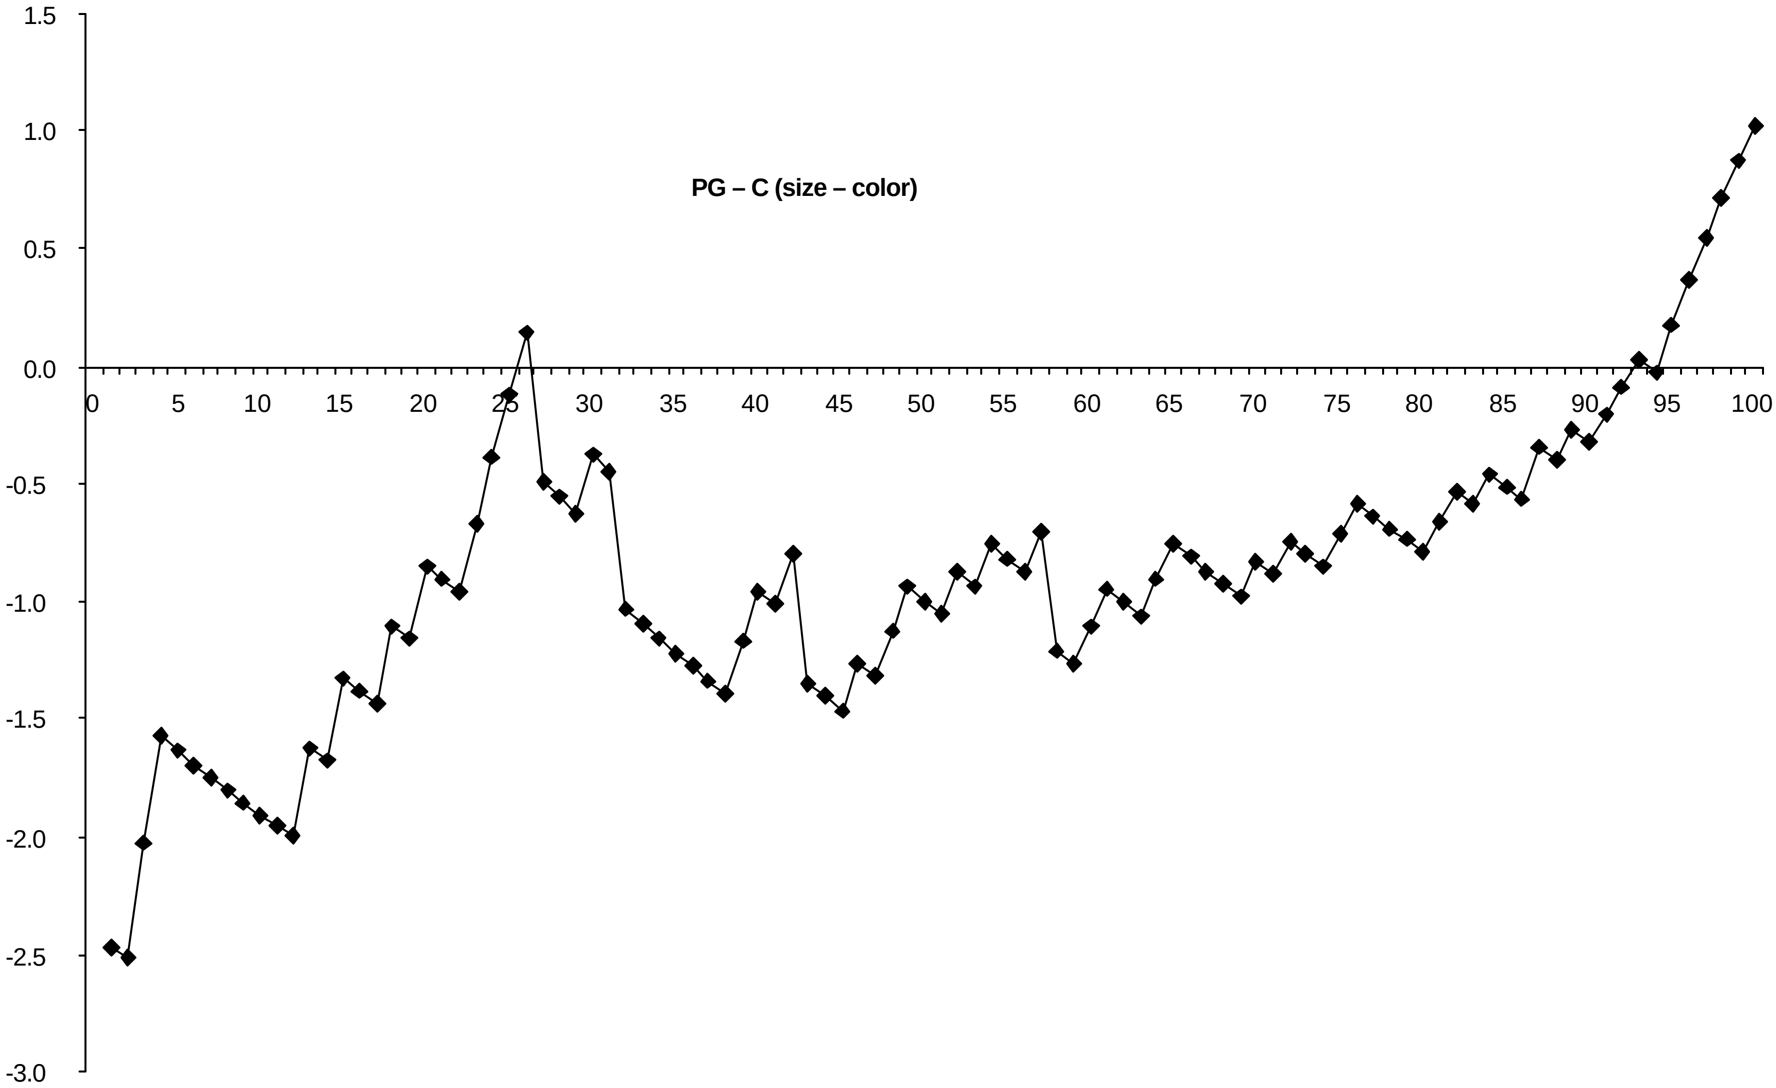

PG–C for Productions *attend-size* and *attend-color*

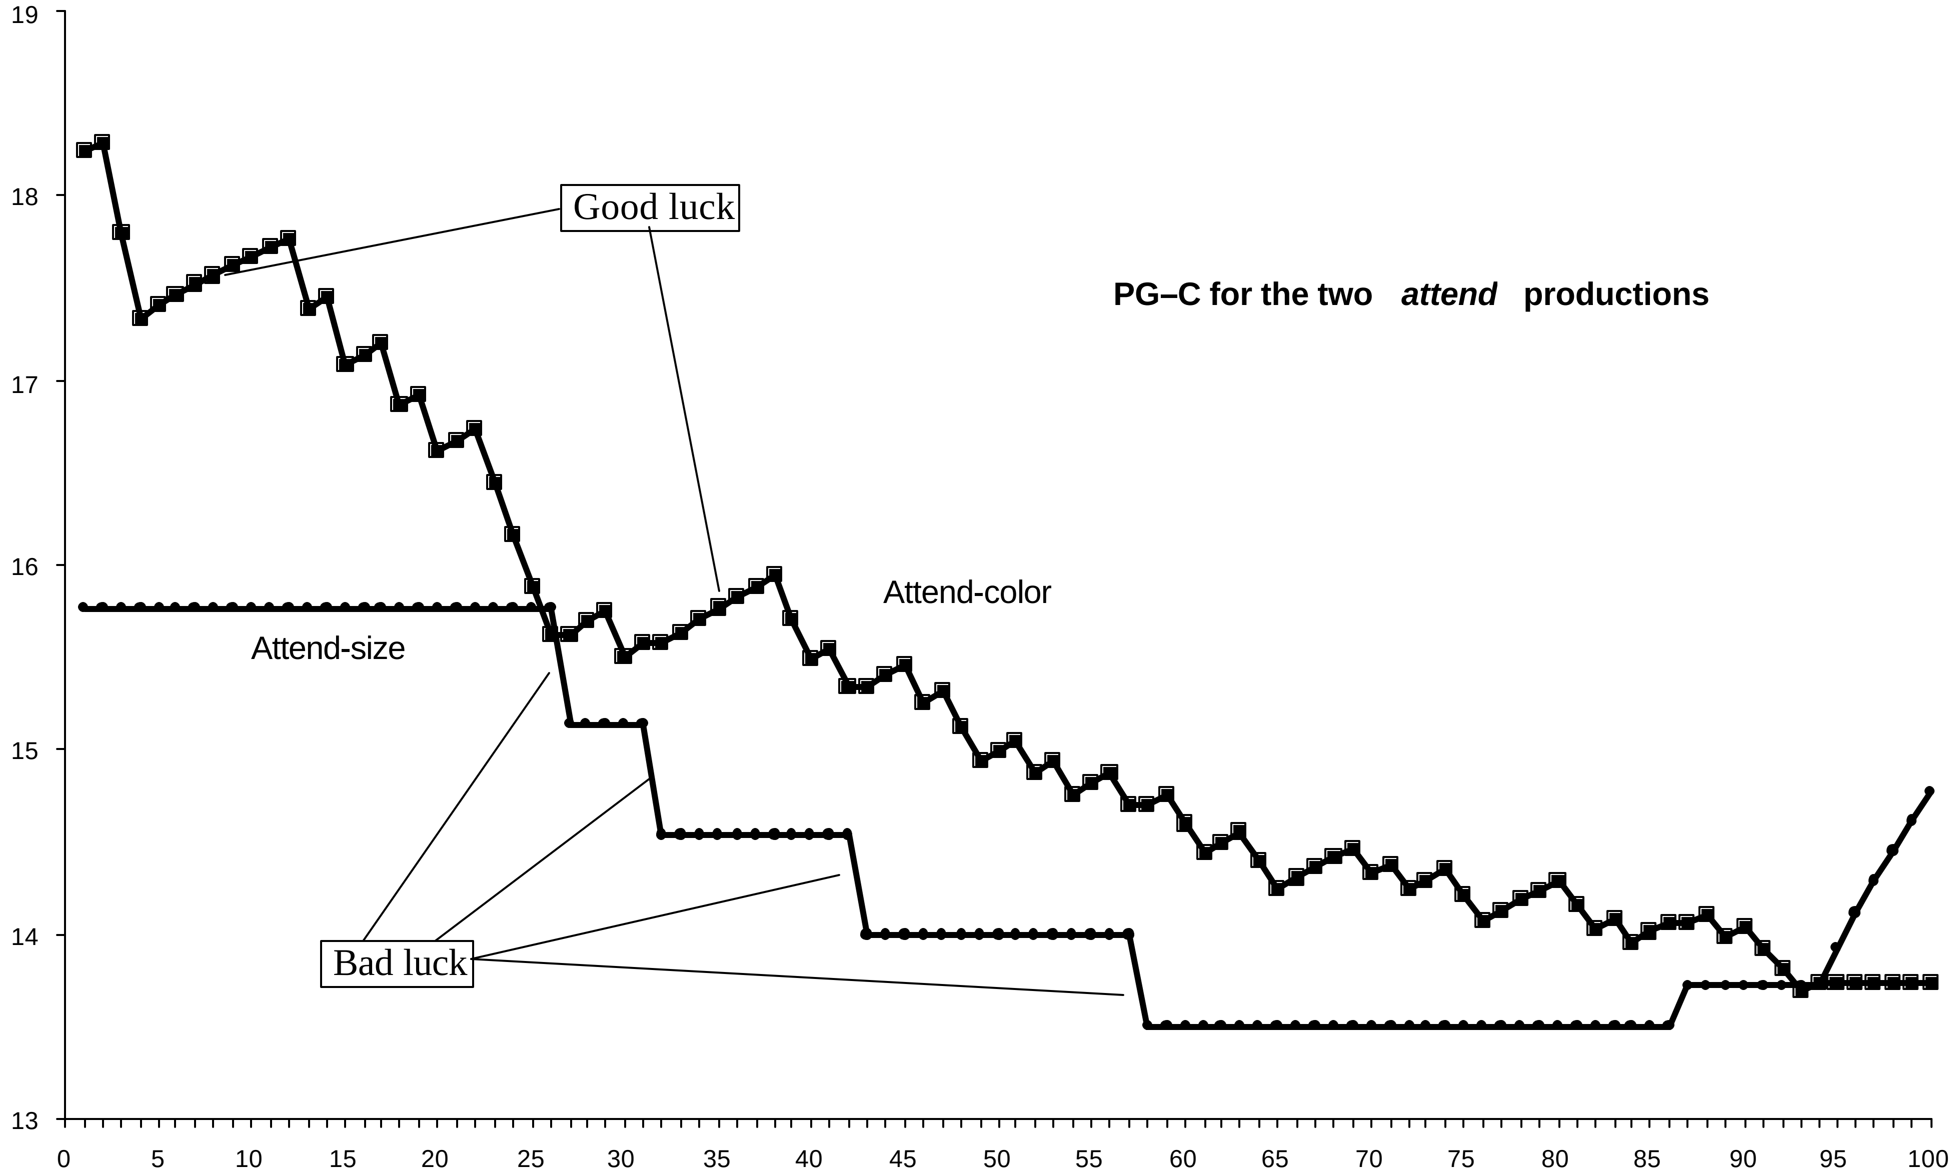

PG–C during Initial Learning

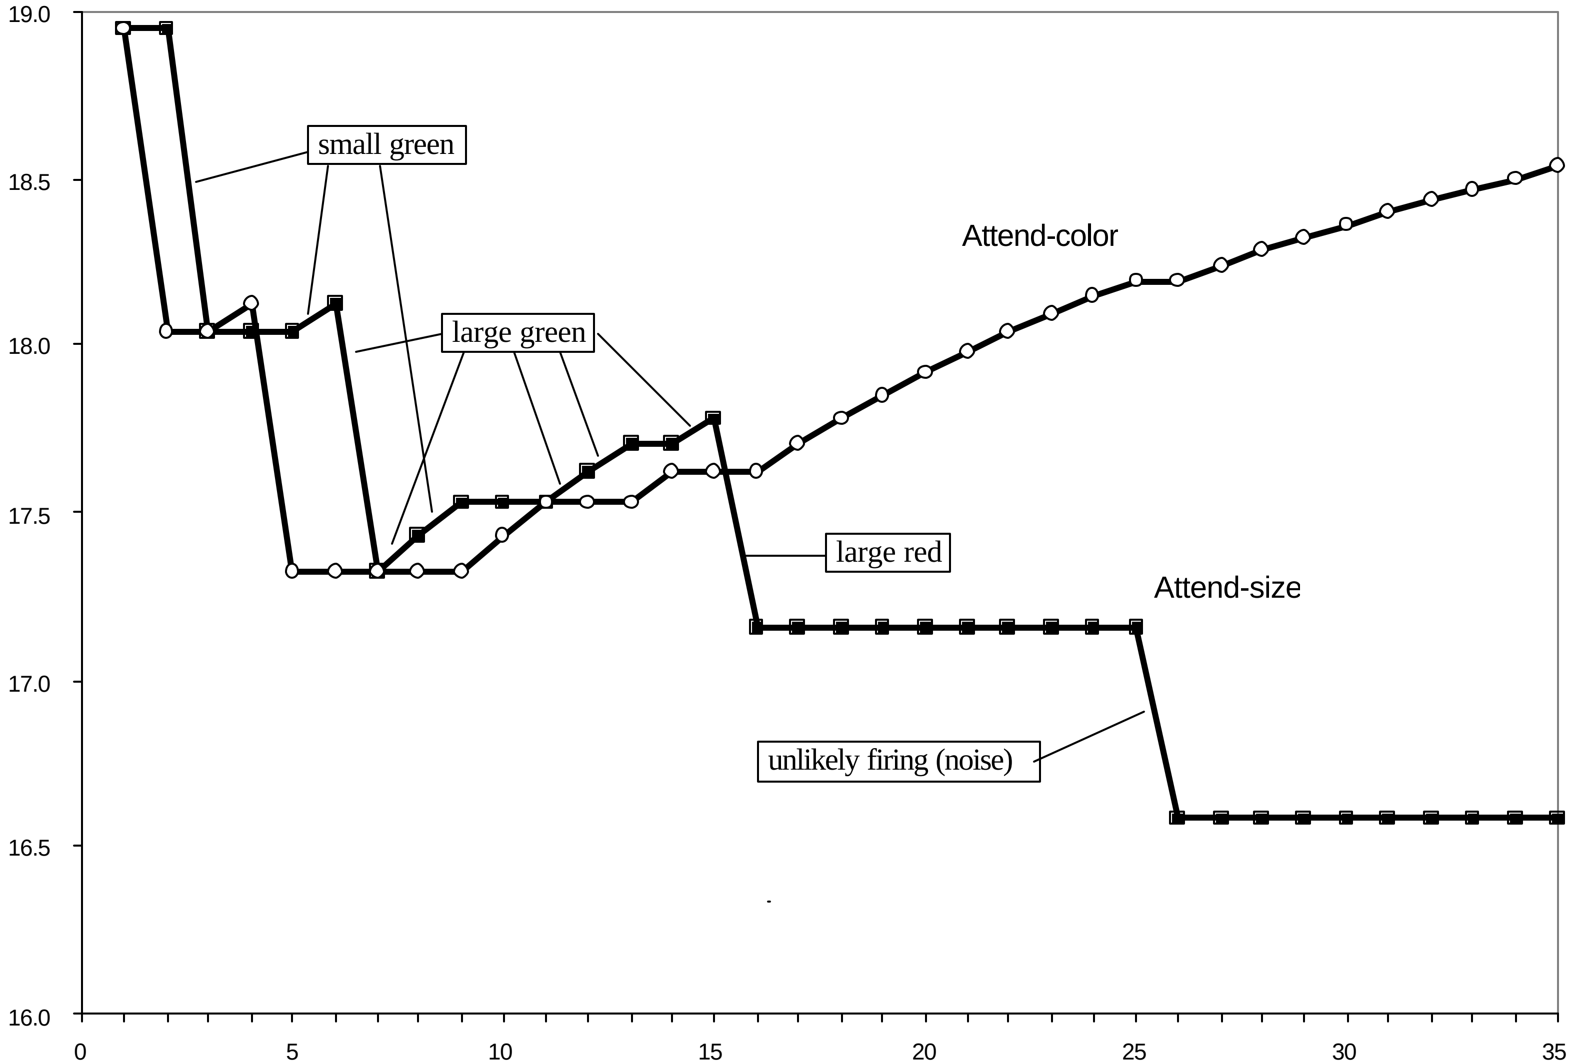

Conclusions

- In Act-R models which use production-parameter learning based upon succeed/fail, and where succeed/fail itself depends on which productions fire, there is an inherent opportunity for random walk processes to occur.
  - — One production can for a while *mask* the other(s), and its PG–C has to random-walk to a lower value before another one can fire.
- These random walks — especially in combination with ‘luck’ — can yield wild distributions of learning times.
- What’s the psychological relevance?
  - — The distribution is a clear prediction from Act-R theory, but do people exhibit that kind of distribution?
  - — It’s therefore a stringent test of Act-R theory.
- A tricky question to answer:
  - — can’t run Ss over and over again, as one can models;
  - — at least some of the variation between Ss is due to stable individual differences rather than random occurrences.
- How to tease them apart?