

Insurance as an ergodicity problem

Ole Peters^{1,2} 

¹London Mathematical Laboratory, London, W6 8RH, UK; and ²Santa Fe Institute, Santa Fe, NM, 87501, USA
Email: o.peters@lml.org.uk

(Received 1 June 2023; accepted 2 June 2023)

In November 2014, the economist Ken Arrow and I had one of our long conversations about my efforts to re-imagine economic science from the perspective of the ergodicity problem (Peters, 2019). Ergodicity, as it pertains to economics, is about two different ways of averaging to deal with randomness. Let us say we measure some quantity at regularly spaced times $t = 1 \dots T$ and model it as a stochastic process, $x(t, \omega)$, where ω denotes the realization of the process. If we want to reduce the process to a single informative number, we can average across the ensemble of realizations, yielding the expected value $\mathbb{E}[x](t) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_i^N x(t, \omega_i)$, or we can average across time, yielding the time average $\mathbb{T}[x](\omega) = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_t^T x(t, \omega_i)$, Fig. 1. If the process is ergodic, then the two ways of averaging will give the same result. We are interested in cases where this is not true.




Figure 1. A stochastic process can be averaged across realizations (top to bottom) or across time (left to right) to produce its expected value or time average, respectively, in the limits of infinitely many realizations or infinite time. If the process is ergodic, the two procedures yield the same result. For many important economic models, this is not the case.

What I had noticed and discussed over some years with Ken was that economics relies heavily on averaging across the ensemble – the expectation operator, $\mathbb{E}[\cdot]$, appears almost universally in papers dealing with randomness, for instance in the context of choice under uncertainty. Averaging across time, on the other hand, is used relatively rarely. But it seemed to me that for individual decision making, there is something unphysical about expected value, and time averages are often a better guide. This is not because of the limit, $N \rightarrow \infty$, implied by the expected value but because my own experience will always be that of one single reality, and even an average over $N = 2$ realizations is literally unrealistic. If I lose in a game of chance, it will be cold comfort




Figure 2. Repeatedly tossing a fair coin for a 50% gain or 40% loss leads to exponentially growing expected wealth (red dashed). This is contrasted with the actual wealth of 100 realizations (orange), which decays exponentially, in the long run at the time-average growth rate (slope of the green dashed line), which is negative in this case: an illustration of ergodicity breaking, where expected value does not indicate what happens over time, developed in detail in Peters and Gell-Mann (2016).

that an imagined “mental copy” of myself, as Schrödinger (1946, p. 3) called other realizations, won the game. I cannot access the ensemble, nothing physical can be exchanged across its constituent parallel worlds. To access what happens later in time, on the other hand, all I have to do is wait and that is (sometimes) physically possible.

The problem can be neatly illustrated with a repeated multiplicative coin toss where heads means a 50% gain in wealth; tails a 40% loss. The expected value grows at 5% per round ($\frac{1.5+0.6}{2} = 1.05$). However, ergodicity is broken so that in the long run each trajectory loses about 5% per round (multiplying the growth factors together gives $\sqrt{1.5 \times 0.6} \approx 0.95$ per round, Fig. 2).

Ken suggested that I should use the insurance puzzle, as he called it, to communicate what we had done. In any insurance contract, the buyer pays a known premium and receives an uncertain payout, contingent on some future event. The fundamental insurance problem, as he saw it, may be phrased like this (Peters & Adamou 2015): No matter how the payout is modeled, the changes in expected wealth for the buyer and seller satisfy the simple symmetry $\Delta \mathbb{E}[x^{\text{buy}}] = -\Delta \mathbb{E}[x^{\text{sell}}]$. Whenever one expects to gain, the other must expect to lose. Since this is the case, why would two parties voluntarily sign such a contract?

Ken was dissatisfied with the most common answers. “Risk aversion” or “utility” is one, but this is either just a relabeling of the observation that one party accepts a poor deal or it refers to someone who systematically behaves sub-optimally, typically an unsophisticated buyer. “Asymmetric information” means that the buyer and seller may have different risk models and therefore both may believe they are getting a “good deal.” But are these good explanations for the existence of the entire industry? If we count the derivatives market as part of it, then the vast majority of all financial activity is of insurance type. Participants in the derivatives market are not unsophisticated and they interact repeatedly – why would they if they frequently found themselves fooled?

The ergodicity solution of the insurance puzzle constitutes a switch of perspective. Since both parties are real entities, not ensembles of mental copies, we analyze insurance contracts as embedded in a trajectory along time, not in a statistical ensemble of trajectories. This means we will compute time-average growth rates (like the slope of the green dashed line in Fig. 2) not changes in expected wealth (red dashed line). For a given model of wealth, finding the functional form of the growth rate, g , whose time average converges meaningfully requires us to take the wealth dynamic into account. In the simple coin toss, the dynamic is purely multiplicative so that wealth grows exponentially, and the appropriate growth rate is logarithmic. Computing its time average for the buyer, $\mathbb{T}[g^{\text{buy}}]$, and for the seller, $\mathbb{T}[g^{\text{sell}}]$, we find that both can gain over time (though not in expectation) by setting up the contract. Insurance becomes a win-win transaction, not one where one party outsmarts the other.




Figure 3. Over time, expected wealth (red dashed) is an unachievable fiction for a solitary agent. In the long run, the 5 agents who judge insurance by time-average growth and often purchase it (blue) exponentially outperform the 5 agents who judge insurance by expected wealth and reject it (orange). The latter lose at the time-average growth rate for uninsured agents (slope of the green dashed line).

This is relevant in practice because real wealth is not well modeled as purely additive. Instead, there really are multiplicative, exponential, aspects of wealth dynamics. A savings account is the simplest example, but also basic expenses on safe housing or nutritious food provide non-additive returns. The same is true of buying a car to reach new places of work or of investing in education. Conversely, a large hit to my wealth, say losing a house to fire, or a car to an accident, has multiplicative knock-on effects. Because expected value is an additive operator, such non-additive dynamic effects break the ergodicity of the linear rate of change of wealth, and thereby the symmetry which creates the insurance puzzle. The change in the time-average growth rate of one party is not the negative of the change in the time-average growth rate of the other party, $\Delta T[g^{\text{buy}}] \neq -\Delta T[g^{\text{sell}}]$. By signing a contract, both parties can boost their time-average growth rates, although it is true that they can never both boost their expected wealth.

To illustrate what this means, Fig. 3 displays the wealths of agents who repeatedly face a 5% probability of losing 95% of their wealth¹. We allow them to insure each other. In the first instance (blue), if agent i faces a risk, he collects insurance offers from the other 4 and signs a contract at the cheapest offer if that increases his time-average growth rate. Conversely, each of the other 4 agents, $j \neq i$, offers competitively priced insurance, namely each at the fee which leaves the time-average growth rate of agent j unchanged. In the second instance (orange), agents act so as to optimize expected wealth and never buy insurance. The overall result is striking: agents who optimize time-average growth frequently buy and sell insurance and thereby exponentially outperform agents who optimize expected wealth.

Pricing individual insurance contracts, from this perspective, is the task of finding a price which constitutes a genuine win-win agreement, where both parties will grow faster as a result. Insurance companies may use estimates of such win-win ranges to identify products for which demand exists. Regulators may use them to determine whether markets are functioning as they should.

If we think abstractly of insurance buyers and sellers, we may ask what sort of entities are likely to be on either side of a deal. Broadly speaking, because a given dollar loss is less disruptive for a large entity, large tends to insure small – a difference in the level of resources is enough to produce a win-win price range. At the systemic level, this raises an interesting question. With premia paid from small to large entities, and most economic activity constituting such transactions, what is the resulting ecology of small and large?

¹Code to reproduce all figures is available at <https://doi.org/10.5281/zenodo.7994190>

With my group at the London Mathematical Laboratory, we are only at the very beginning of exploring these systemic dynamics, and what we find continues to surprise us with its richness. If everyone signs insurance contracts when that is in the agent's long-term interest, the richest agents will stay on top for long periods. However, they can never fully break away from the herd. As they grow, they eventually run out of willing insurers and have to fend for themselves. This can lead to sudden rearrangements of the wealth hierarchy. The rich lose because they lack partners to cooperate with.

We can view insurance as a lifestyle product, allowing us not to worry about scratches in the holiday rental car; we can also view it as a product for the risk averse, who do not like ups and downs. Ergodicity economics exposes a view of insurance as a tool for boosting growth (or slowing decline). Risk management, here, is not merely about reducing ups and downs but in the long run about creating the kind of stable conditions which nurture and boost multiplicative growth wherever it occurs.

Ken Arrow, sadly, passed away in 2017. I feel that his advice was sound: pairing ergodicity economics with actuarial science creates a large win-win space which I invite readers to explore.

Funding statement. The ergodicity economics program at the London Mathematical Laboratory is supported by Baillie Gifford and by the Novo Nordisk Foundation (Exploratory Interdisciplinary Synergy Grant, ref NNF20OC0064869).

Competing interests. None.

References

- Peters, O. (2019). The ergodicity problem in economics. *Nature Physics*, **15**(12), 1216–1221. doi: [10.1038/s41567-019-0732-0](https://doi.org/10.1038/s41567-019-0732-0).
 Peters, O. & Adamou, A. (2015). Insurance makes wealth grow faster. <http://arxiv.org/abs/1507.04655>.
 Peters, O. & Gell-Mann, M. (2016). Evaluating gambles using dynamics. *Chaos*, **26**(2), 23103. doi: [10.1063/1.4940236](https://doi.org/10.1063/1.4940236).
 Schrödinger, E. (1946). *Statistical Thermodynamics*. Cambridge, UK: Cambridge University Press.