

Luca Dellanna on Why Long-Term Investors Must Understand Ergodicity

We had the pleasure of speaking with Luca Dellanna, author of [Ergodicity: How Irreversible Outcomes Affect Long-Term Performance](#).

We are grateful to Guy Spier for the introduction to Luca.

This conversation is available as an episode of Invest Intelligently, a member podcast of MOI Global. (Learn [how to access](#) member podcasts.)

See also [Alex Gilchrist's conversation with Luca](#) on the same subject.

The following transcript has been edited for space and clarity.

John Mihaljevic: It is a great pleasure to welcome to this conversation researcher, author, and management advisor Luca Dellanna. I've known of Luca for quite some time and have read a couple of his books, but I must give credit to my friend Guy Spier, who I believe is the first person to bring Luca to my attention. He also helped make this interview happen, so thanks for that, Guy.

Luca, I know you're an automotive engineer by training and have led teams and consulted for large multinationals. Then you essentially quit your job to do what you do now. I think this is always admirable and exciting, especially when you bring so much value to all of us. As a start, I'd love it if you tell us a little about what you are up to now these days, how you spend your time, and what you focus on.

Luca Dellanna: Thank you, John. Nowadays I spend half of my time writing and researching, usually around the topics of risk, people management, and operations management. The other half of my time I spend consulting. I usually work with senior leaders to improve the management capabilities of their organization, so I do things like training the managers and training the leaders in developing the managers.

John Mihaljevic: Today I'd love to talk to you about one of your books, the one on ergodicity. In a future conversation, maybe we can also talk about your latest book, which is called *Winning Long-Term Games* and also looks extremely interesting. The subtitle in *Ergodicity* is "How irreversible outcomes affect long-term performance in work, investing, relationships, sport, and beyond." It's a fascinating concept. We in the MOI Global community are investors and lifelong learners, so this is extremely relevant. As a starting point, how do you define ergodicity? How should we think about this?

Luca Dellanna: I think the best way to introduce ergodicity is not to talk about its definition. I use an example in the book that really brings the idea across. It's a story of my cousin, who was a great skier. He made it to the World Championship in his age bracket when he was very young, but with one leg injury after another, he sadly had to quit professional skiing very early. From him, I learned the lesson that it is not the fastest skier who wins the race but the fastest one amongst those who make it to the finish line.

In the book, I use a numerical example to explain the point that not only does survival matter (we all know this) but that survival matters more than performance, at least over long-time frames. Here's the example I make. Imagine a ski championship consisting of ten races. My cousin is a great skier, so he has a 20% chance of winning each race. However, because he takes a lot of risks, he also has a 20%

chance of breaking his leg in each race. The question is how many races he is expected to win in a championship of ten races.

The naive answer is two races because we think, "There are ten races and a 20% chance of winning each; ten times 20% makes two." However, the real number is only 0.71, because if my cousin breaks his leg during the first race, not only does he lose that race, but he also loses the possibility of participating in future races. He only has an 80% chance of participating in the second race, a 64% chance of participating in the third race and so on.

We get this huge difference. In a world where irreversible losses do not have long-term consequences, my cousin wins two races. In the real world, where irreversible losses absorb future gains, my cousin only wins 0.71 races. This difference caused by irreversibility is what ergodicity is about. A situation is ergodic if irreversible losses do not have long-term consequences and you can rely on averages. The real world is non-ergodic, which means that 1) irreversible losses absorb future gains and 2) because of that, you cannot rely on averages.

John Mihaljevic: That's a great example. To stay with the skiing analogy for a moment, there are some great skiers or have been in the past who end up winning many of the individual races they enter, and they also don't get taken out of the race. What does that say about their skill level? Are they not skiing at 100% and are they just so much better than the rest that they take fewer risks and still win individual races?

Luca Dellanna: Part of the answer to this is survivorship bias. We are seeing the subset of skiers who are running at 100% and survived, but we are not seeing the subset of skiers who didn't survive, such as my cousin. This is one part of the answer. The second part of the answer is that skiing is a race to the bottom, by which I mean that in professional skiing, you can only have one winner. This means two things. If you are very skilled but also don't take many risks, someone almost as skilled as you but taking more risks will be the one winning. Of course, there will also be a lot of other skiers who are just as skilled, take a lot of risks, and get out of the race because they injure themselves. We do not see them because of survivorship bias.

The point is that since there can only be one winner in professional skiing, there is this race to the bottom where you do not win only because you're the most skilled person but also because you take risks beyond the level which would be optimal if there could be more winners. Conversely, in areas such as investing, there can be more than one winner. Of course, only one gets to be the richest person in the world, but if you define winning as getting very wealthy, as multiplying your wealth, there can be more winners, which means you do not need to take a suboptimal level of risks to be the winner. There are pathways to being a winner where you take a sustainable amount of risk.

The good way to think about it is to get rid of survivorship bias. You want to look at winners and ask yourself, "If we take 100 people with the same strategy as theirs, how many of them ends up winning?" If the answer is not 100, then you have a strategy which is not reproducible, meaning that if you were to adopt the same strategy, you are not guaranteed those results. Non-reproducible strategies work when you have a lot of people adopting them so that some drop out of the race but there are still some winners. They do not work when you are thinking about whether they are a good strategy for you because you are only one person and if you blow up, you blow up. You want to optimize for reproducibility before you optimize for maximum potential.

John Mihaljevic: Would it be fair to say it takes luck to win in skiing? It almost sounds like it does take luck, but more to not have a bad injury, to not have the irreversible loss. Let's say the top 10 skiers in the world are pretty much the same in terms of skill and, as you say, it's a race to the bottom. They all push it almost to the max. It's almost like the one who was the luckiest in not getting injured

ends up the overall winner.

Luca Dellanna: It's a big component you're mentioning. There is often this belief that if we say luck is important, then it means that skill is not important. It's not like that. Skill is extremely important, and so is hard work. Everyone who wins a race is not just lucky but also incredibly skilled, and they put in an incredible amount of hard work. It's just that when you can only have a limited number of winners, skill and hard work are not sufficient anymore, and you also need to have a bit of luck.

John Mihaljevic: Would you say that the skier you mentioned at the outset was, in a way, rational in what he did because maybe that was his only way to have a chance at winning, and he took that risk of losing everything? That's just how the game works. If you want to have a chance at winning, you also have to take that risk.

Luca Dellanna: You're highlighting a very important point. The moment you decide to compete to win in a race to the bottom, if your objective is to become number one, then the rational thing to do is take a high level of risk. What's irrational here is the decision to compete in a race to the bottom in the first place. It's my opinion. That's the big topic in the other book, *Winning Long-Term Games* – you should not compete in races to the bottom to win because those are races where to win, they make you do things with a high potential of being bad for you.

The rational thing is to not participate in a race to the bottom. If you're competing, try to find competitions which can have a large number of winners. For example, you compete in investing, but you do not compete to become the wealthiest investor in the world. You just compete to grow your wealth a lot, which is a type of competition that doesn't mislead you because the rational thing to do is not to take a suboptimal amount of risk anymore. My advice is to make sure you pick competitions in which you do not have to take an excessive amount of risk in order to have a chance of winning.

John Mihaljevic: Turning to investing, where anyone can be a winner based on that definition, I'd love to understand a little more the nuances of how you look at the concept of irreversible losses absorbing future gains. Mathematically speaking, is it the idea of a multiplicative series having one zero? It's all a zero or a kind of, if you're down 50%, you need to be up 100% to get back to even, and if you're down more, you need to be up a lot more. How should we think about this mathematically?

Luca Dellanna: It's both things you're describing. If you go bankrupt, are forced to liquidate an investment at an awful price or have so many losses that your investors decide to cash out, this basically brings you to a zero, a failure point (the technical term is "absorption barrier"). The other phenomenon is multiplicative risk: if you go up 50% and then go down 40%, you have actually lost money. You are not up 10% but below where you were at the beginning. That is also part of it. One way to see it is that if you invest \$1,000 and then lose 50% of your investment, you did not just lose \$500 but also the future profits those \$500 might have generated.

John Mihaljevic: Looking at some of the most successful investors, like Warren Buffett, it sounds like they do grasp this concept very well intuitively. Do you think the longevity of a track record has a certain quality to it itself, and not just the returns because people will often compare returns? It sounds like the longevity of those returns is maybe even more important.

Luca Dellanna: Exactly. Because of compounding, longevity is extremely important. We all know Warren Buffett made most of his wealth during the last years, and that's not because he had a higher rate of return but because of the effect of compounding. You do need the longevity because if you look at absolute numbers, the big gains are made later in time. To get into longevity, you need to have the ability to survive.

The longer your returns are sustained, the more likely it is that your strategy will produce future returns. Nassim Taleb calls it the Lindy Effect in his book *Antifragile*. It is the idea that the longer something has been around, the longer it is expected to survive, but it does not apply to things that have a limit to how long they can survive, such as humans. If you survive 90 years, you are maybe expected to survive another five years, but there is a hard limit. For things that do not have this natural bound — such as ideas, investments, and technologies — what you observe is that the longer they have been around, the longer they are expected to be around in the future.

The underlying reason for the Lindy Effect is that if you survive for long, we can deduce that your hazard rate, the chances of dying every any single year or your chances of going to zero in any single year, are low. The longer you survive so far, the more certainty we have that your hazard rate is low, and therefore the more likely you are to survive in the future. For example, for Warren Buffett, who sustained the long returns for many decades, we can be rather confident that his hazard rate is low, whereas if you pick a random investor who has very high returns but has been in the game for only the last two years, we do not know their hazard rate. The hazard rate might be high, but the investor has not gone bankrupt because he got lucky in those two years. Two years are not enough to close the cycle, so maybe the investor only experienced good market conditions.

So, there is this aspect of longevity. One, longevity is something you want to look for yourself. Two, how long an investor has been around gives you some information about how reproducible their strategy is. With a strategy that looks successful but has only been around for a while, you do not know yet whether it is reproducible, which means if you were to adopt it, you do not have a high degree of certainty that you will get the same results.

John Mihaljevic: That is a great point. Related to that, I feel like the investment industry has invented things like risk-adjusted returns and beta, which, to me, look more like marketing tools to say, “Even though I haven’t been around long, my returns haven’t been that risky” based on whatever metric they want to use. It seems like there’s nothing as good as simply having a long-term track record. With Warren Buffett, I guess you could say his reported returns are the risk-adjusted returns because he’s been around so long that if there had been some major risks, he would have blown up by now. How do you view some of these risk-adjusted return measures, and do you think it is possible to try to ascertain with some confidence, even when there is not a long track record available?

Luca Dellanna: I largely share your thoughts about metrics, which I think are more like marketing tools or tools to increase the confidence in investors so that they’re more likely to invest or assume some risk they will not assume otherwise. I personally dislike using most metrics because they might fit the past but don’t tell you much about how they will fit the future. A lot is not considered in many metrics.

Take, for example, two investment portfolios, both with an average rate of return of, say, 20%, which is very good. They both have some metrics about risk and other things, but the first portfolio has been around for two years and second for 20 years. To me, the second portfolio has a much lower risk because of the fact that it survived for so long, which means it’s also more immune to things such as survivorship bias. However, that’s not captured by many metrics. I think people should never rely excessively on metrics. I often advise my clients to rely more on thought experiments, like not thinking, “What are the chances I have a big loss?” but rather think, “Let’s assume my portfolio has a very big loss. What’s the most likely reason?” Then you can make some considerations about that.

Another thought experiment I like is to think, “If 100 people take this portfolio, what’s the distribution of outcomes for those 100 people?” You’re shifting from thinking in terms of averages and thinking in terms of potential to thinking in terms of distribution of outcomes, which makes the bottom part of the tail much more visible. That’s the way I prefer to think about investments – in terms of thought

experiments, in terms of premortems. You do not ask yourself how likely the possibility of failure is. You assume failure and ask yourself what the most likely cause could be and then you make some considerations about how likely it is to happen and what you can do today to prevent it.

John Mihaljevic: Some folks mentioned Monte Carlo simulations. What role could those play?

Luca Dellanna: I love Monte Carlo simulations as a tool. I think that's very effective, for example, to compute the distribution of outcomes. It helps you capture things like irreversibility and ergodicity. The limitation of Monte Carlo tools is that they are only as reliable as the variables you put inside them. If you know all the variables and the uncertainty and are correct with that, then the Monte Carlo tool will be extremely effective. However, if you have high uncertainty on the variable you're putting in, or at least high relative uncertainty, then Monte Carlo might be misleading. It can help you detect some blind spots, but it will not help you detect all the blind spots in your strategy.

It's a helpful tool, but it's not a tool you can blindly rely on, and sadly, most people underestimate how much uncertainty they have on their variables. For example, people think, "I know that variable is likely to be in the 2% to 4% range," but that's not the thing that matters. The thing that matters is what happens in the 5% of cases where it's outside of that range. Is it, for example, 5% or 10%? In the 5% of cases where it's outside the 2% to 4% interval, is it 1% or is it minus 10%? That's what you really should think about.

John Mihaljevic: Based on your research, have you come to some practical conclusions about portfolio management in terms of how diversified a portfolio should be, how much leverage is "safe," and anything else investors can apply when it comes to constructing their own portfolios?

Luca Dellanna: I'm not a professional when it comes to investment advice, so take what I'm saying with a grain of salt. My suggestion will be to think, first of all, what your real time horizon is. You may say your time horizon is 10 years, but unless you are 90 years old and your life expectancy is 10 years, your time horizon is not 10 years. Your time horizon is the top of the 90% confidence interval on your conditional life expectancy. For me and you, our real time horizon is probably something like 50 years.

Now the question is what's the maximum amount of risk you can rationally take with such a long-time horizon. You may consider a strategy with a 1% chance of blowing up a rational one for a 10-year time horizon portfolio because maybe you're thinking, "Yes, there is this little risk of blowing up, but the gains are so high." Over 50 years, that's definitely not worth it. The thing you want to consider is if you want to have a 99.9% chance of having a good portfolio after 50 years, which is your real time horizon, then the real amount of risk of blowing up you can take any single year is tiny, tiny, tiny. That's one consideration. By the way, blowing up doesn't necessarily mean you lose 100% - maybe it's something like 90% or 80%, which could still qualify as blowing up.

The second consideration is, let's say using this thought experiment, you discover that the yearly risk you can take of considerable losses is one in 10,000. Again, this seems too tiny, but if you're there for the next 50 years, it's not necessarily too tiny. Once you have a strategy, you need to think, "Let's imagine 10,000 parallel worlds, and I need to have a positive outcome in 9,999 of those 10,000 worlds." Then you ask yourself if your current strategy fits this criterion.

Those numbers are maybe a bit extreme. Some investors will have different numbers. Maybe you discover your acceptable risk is only 1% or whatever; the thought experiment is still valid. You ask yourself, "In how many parallel worlds do I get a positive outcome that I'm satisfied with?" and this should be, in my opinion, one of the compasses guiding your decisions. Once you're okay with it, there are plenty of strategies you can consider. The point is to make your considerations for your real time horizon, which is your life expectancy, not 10 or 20 years. The second thing is to think in terms of

parallel worlds, and you want to have a positive outcome in all of your parallel worlds.

John Mihaljevic: That's very helpful, thank you. Anything else in the book you would like to highlight on the topic of ergodicity?

Luca Dellanna: One thing I talk about is the importance of not concentrating your risks over the time axis. I use an example that has nothing to do with investing, but I think it's useful to understand the concept. I imagine I'm going on holiday with my wife, and we've decided to drive there. It's a very long car drive, say, 12 hours, and if I drive without stopping for the night, I'm causing problems to cluster towards the end of the drive because then my wife and I will be the most tired. Maybe we will be the most frustrated, and the likelihood is that if something goes wrong in that last part of the drive, we'll feel it very much. Because we are both tired and a bit frustrated, that makes it more likely that we have an argument, and if we have an argument while we are tired and frustrated, it makes it more likely to be a bad argument.

One simple way in which I can reduce this risk is to break the drive in two parts, have an overnight stay so that there is no point at which we are the most tired at the same time, which is a factor that might create a problem. This concept is also useful because we often think we have a certain capacity to make decision but the problem is our decision-making is usually better when we need it the least and it's the worst when we are stressed because we have a situation that requires decision-making. Another problem, for example, is that usually our investors trust us the least the moment when we need their trust the most. Risks and problems tend to cluster in the same block of time, making risk multiply in a way that is not apparent in simulations or many risk models.

One thing you want to think about is how you can make sure that those risks are as distributed as possible over time. You probably know the specific tactics better than me. I know, for example, that some investors only require clients to make decisions like whether to stay invested or not once every quarter. I know that some investors try to focus on ways to spread the risk over time. If they know that certain periods will have concentrated risk, they take precautions in advance. So, think about this aspect, that risks tend to cluster together, and think how to decluster risks so that they do not happen all the same time. If you can do that, your decision-making will be the best when it's required and your clients will be willing to give you trust when you need it.

John Mihaljevic: I guess the concept also applies over a lifetime when I think of one's reputation, how we can use ergodicity to illuminate what I think most of us understand intuitively in terms of how we treat other people or cut corners in life.

Luca Dellanna: Yes, that's an excellent point, John. Reputation is also very much non-ergodic in the sense that you do not really care about the average reputation of a person's actions because it only takes one bad action to destroy the trust you have in someone or the trust people have in you. The problem is that this is tricky, and I use the example of lying in *Winning Long-Term Games*. Let's imagine lying only has a 1% chance of getting you caught. This makes it seem a great strategy because it means that by lying, you have a 99% chance of getting away. However, if you lie every week, you have a 99.5% chance of getting caught over five years.

The reason I'm talking about this is because there are a few investors – and a few people in general – who take actions that put at risk their trust, but because the chances of destroying trust are small, maybe they get away with it for a few years, which makes them think it's a great long-term strategy. However, it's not because over the long term, you will get caught, your reputation will be destroyed, and that will be terrible.

Here's where it gets tricky. Let's say you truly believe that you should always be sincere with your

clients, and you start working like that for one year, two years, three years. Then you notice there are a few competitors of yours who instead lie to their clients or maybe exaggerate some potential returns. Because of that, they are becoming more successful than you. At this point, you feel like you're falling behind, like maybe their strategy is better. It takes a lot of determination to avoid switching from your good, long-term strategy to the worse strategy of these other people, which only looks better because we are examining it over the short term.

I always give the example of the casino croupier. In the casino, the croupier is the only person at the gambling table who has a money-making strategy. All the players have a strategy that is expected to lose money over time. However, every day, the croupier sees at least one player getting wildly rich, and they must resist the temptation to switch from their good, long-term strategy to the worse strategy of the players, which is a money-losing strategy even if today it looks like it produced a lot of returns.

The same applies when we look at our competitors. Perhaps they have a strategy that looks better in the short term, maybe a strategy of lying or exaggerating returns, but it looks better only because we're seeing the short-term consequences. In the long term, it's a worse strategy, but it takes resolution and long-term evaluations to have the determination to stay with our good long-term strategy, even when it feels like we're falling behind in the short term.

John Mihaljevic: If we wanted to be a little cynical, we can say that a lot of clients or people in general tend to have short memories. A country that defaulted on its debt 10 years ago can now raise debt again from investors who weren't involved before and don't have that memory. Similarly, a hedge fund manager who blew up once goes out all of a sudden some years later and is able to raise a lot of money again. How can we work the short memory issue into this analysis? How do you think about that?

Luca Dellanna: There are two parts to the answer. First, from a societal point of view, I think journalism and similar professions should do a much better job of reminding people of what happened in the past and of the risks that are still there even though they haven't manifested yet because of survivorship bias and luck. The second part is that investment funds and similar entities have to educate clients, make them understand the risk, make them understand that two percentages that look the same on the paper of return are not the same if they were achieved taking different risks. The risks are not always quantifiable because a lot of quantifiable looks at what manifested in the past, not necessarily what manifests in the future.

These are things people should be educated about. I know some people don't have the language or the examples to educate, which is the main reason I wrote *Ergodicity* and *Winning Long-Term Games*. They contain language and examples people can use to educate others. The example of the skier and other examples in the book can be used when you're talking to clients and explaining why this thing that looks successful is actually risky or why you're aiming for a return that's a bit lower than what your competitors are aiming for (because you can achieve it with much higher reproducibility). There are examples and stories you can use, and they go a long way in educating clients.

John Mihaljevic: I'm really curious about something, so let me ask you this final question. You are so productive and prolific that I'd love to hear about your writing process and how you manage to bring so much value to all of us.

Luca Dellanna: Thank you. I think my key to writing well is that I only write about things I've worked on in the past. All of my books come from my consulting work. When I notice that my clients have a recurring problem or a recurring question and there is a recurring solution that works, I know I have material for a book.

In my opinion, this is a good approach for two reasons. One, it's more likely that the book's contents are useful to the readers. Two, when you write about something you've discussed over and over with other people, writing becomes so much faster. This is why I'm able to produce books so fast - the writing process started years earlier with all the conversations I had with my clients, and that's incredibly helpful. Some try to write in a vacuum. They don't start with "I have something that would make a good book" but with "I want to write a book. What can I write it about?" This makes it less likely that the book will be about something useful. It also makes the writing process much more painful because now they need to think of what to write about and how instead of simply referring to conversations from the past.

John Mihaljevic: That's a great approach. Luca, thank you so much for taking the time to have this conversation. I learned a lot, and I'm sure our members did as well. I hope to reconnect with you on another occasion to talk about one of your other great books.

Luca Dellanna: Thank you, John, it was my pleasure. I would also love to chat with you again.

About the interviewee:

A mechanical engineer by training, Luca Dellanna decided to quit his corporate job to become an independent researcher and author and shed light on the topics of nonlinearities on human collective behavior. Luca believes that those topics are essential for preventing human suffering, especially as the scale of our civilization keeps increasing. His style is concise and direct, focused on cause-effect relationships. He rejects top-down theories and explains most real-world phenomena with bottom-up hypotheses. Luca published his first book, "The Control Heuristic: Explaining Irrational Behavior and Resistance to Change", in 2017. In the following year, he published another book titled "The World Through a Magnifying Glass: A New Theory to Explain Autism" and followed it up with "The Power of Adaptation". Luca writes regularly on Twitter (@DellAnnaLuca). His personal website is luca-dellanna.com.