## Journal of Cognition and

## The Enigma Of Probability

Nick Ergodos

## Biography

I have studied logic, mathematics, physics, and philosophy. My academic interests are in the cross section between mathematics, physics and philosophy. A future project is to investigate the relationship between the probability concept proposed here and the logic of Quantum Mechanics in more detail.

## Acknowledgements

I want to thank the unnamed referee of the Journal of Cognition and Neuroethics for many helpful suggestions. I also want to thank Zea Miller for proofreading my text.

## Publication Details

Journal of Cognition and Neuroethics (ISSN: 2166-5087). March, 2014. Volume 2, Issue 1.

## Citation

Ergodos, Nick. 2014. "The Enigma Of Probability." Journal of Cognition and Neuroethics 2 (1): 37-71.

# The Enigma Of Probability 

Nick Ergodos

> I can stand brute force, but brute reason is quite unbearable. There is something unfair about its use. It is hitting below the intellect.

- Oscar Wilde


#### Abstract

Using "brute reason" I will show why there can be only one valid interpretation of probability. The valid interpretation turns out to be a further refinement of Popper's Propensity interpretation of probability. Via some famous probability puzzles and new thought experiments I will show how all other interpretations of probability fail, in particular the Bayesian interpretations, while these puzzles do not present any difficulties for the interpretation proposed here. In addition, the new interpretation casts doubt on some concepts often taken as basic and unproblematic, like rationality, utility and expectation. This in turn has implications for decision theory, economic theory and the philosophy of physics.


## Keywords

Bayesianism, decision theory, expectation, expected utility, expected value, fair game, measurement problem, probability interpretation, St Petersburg paradox, two-envelope problem

## Historical Introduction

Today we have a whole zoo of probability interpretations. We have propensity interpretations, frequency interpretations, objective Bayesian interpretations, subjective Bayesian interpretations, logical interpretations, personalistic interpretations, classical interpretations, formalist interpretations and so on, almost without end. These interpretations claim that the ontology of probability is either physical, psychological, epistemic, logical or mathematical. Some of them view probability as objective, others as subjective. Some even go to the extreme and say that probability is merely an empty word that we are allowed to interpret in any way we want, as long as we do not violate the axioms of probability.

To understand why we have this wild bouquet of philosophical interpretations it is necessary to study history. The cause of the confusion is an old gambling problem called the St. Petersburg paradox. As this problem is still unsolved, it continues to infuse
confusion. The various ways that have been proposed to escape this problem have led to the scattered philosophical situation we have today for the probability concept.

It is easy to state the St. Petersburg problem. Let us say we play a very simple game where you toss an ordinary coin until heads comes up. If heads comes up in the first toss you will get one dollar from me. If it comes up at the second toss you will receive two dollars from me and so on. We double the amount for each tails-toss you manage to get before you get heads and the game ends. The question now is what you would be willing to pay to play this game. The smallest amount you can win is one dollar so the game should at least be worth one dollar. But exactly how much more than one dollar?

The classical answer to these questions is to calculate the expected value of the game and make sure you do not pay more than that. The problem with this approach is that the expected value of this game is infinite. This means that whatever I demand you to pay for the privilege to play this game you should accept the offer. This is because any amount, no matter how big, is smaller than infinity. But this advice is just crazy. No one in her right mind would pay even a modest sum for this game. Something must be seriously wrong here, but what? This is the St. Petersburg problem.

There has been three main ways to attack this problem (Dutka 1988; Jorland 1987).
(a) The advice to pay infinitely much for the game is actually correct in theory but in practice it is not. If you are lucky you could win more money than I could possibly pay you, and even more money than all the money in the world. Likewise, there is a possibility that heads never comes up and we run out of time, because none of us have unlimited time at our disposal. As there are limited resources of time and money in the world the game as stated cannot actually be played in the real world. There is no need to modify the theory - we only have to keep in mind the actual circumstances when it is used. The theory in combination with the actual physical constraints at hand will produce a correct, finite, result. I will call this the Finite World argument.
(b) The advice is mathematically correct but human beings do not value money linearly as the mathematical theory implicitly assumes. What we need is a new theory that complement the mathematical theory whenever humans and money is involved. I will call this the Human Value argument.
(c) The advice is not correct and the mathematical theory therefore needs to be changed or re-interpreted. A re-interpreted theory usually does not give any advice for actions at all. At least not for single cases. The theory might give some
advices for action if a large number of repetitions of a game is considered. Very few attempts have been made to change the mathematical theory itself. I am only aware of one attempt, which we will study separately. I will therefore call this the Reinterpretation argument.

That a simple game like this can lead to such a big discussion was disturbing. And it got even more disturbing as the years, decades and centuries passed without no one being able to solve it. Early on people drew the conclusion that the concept of expected values does not seem to be as natural and unproblematic as was first assumed. The concept of probability, however, is confined to have a value between zero and one, and can never be infinite. This makes this concept immune from ending up in infinity-paradoxes like this. It must therefore be safer to have at the core of the theory. Consequently, soon after the discovery of the St. Petersburg problem the concept of expected value was replaced by the concept of probability as the central concept of the theory. The theory that up to now had been called many things, but never something including "probability," started to be called the Theory of Probability by everyone. The earlier focus in the theory on how to bet on different games of chance was gradually replaced by a focus on pure probability problems.

This was a smart move. However, the concept of probability and the concept of expected value are mathematically very close to each other. If one of these two concepts has philosophical problems connected to it, the other one will have philosophical problems as well. Not exactly the same, as we will see, but similar. The different proposals on how to get rid of the St. Petersburg problem is the direct cause of why we have the probability interpretations that we have, and why they are designed as they are. The Finite World and Human Value arguments are connected to the subjective, personalistic and Bayesian interpretations of probability. The Reinterpretation arguments led to the frequentist, propensity and formalist views of probability. Incidentally, in recent decades the Human Value argument has also played a crucial role in the development of economic theory, game theory, decision theory, and rational choice theory (Samuelson 1977).

The contemporary way to view the St. Petersburg paradox is that it is a very important historical problem that has led to a number of important theories. The fact that the problem is still unsolved does not seem to bother anyone anymore. In fact, the common understanding today is that the problem really is solved, only that it is not possible to say which solution is the correct one... We are told that we are free to pick any of the proposed solutions and let it be the solution of our choice. Notice that the three solution strategies above contradict each other. They cannot all be correct at the same
time. Being a simple mathematical problem this is really an odd situation. In no other area of mathematics are we free to choose the solution to a problem ourselves, and whichever solution in a set of mutually contradicting solutions we pick, we will have picked a correct one. Incidentally, the solution we pick also, to a large extent, determines the probability interpretation of our choice, and vice versa.

However, it is relatively easy to see that the Emperor is naked, i.e., that none of the proposed solutions is a valid solution. By doing this all the theories, concepts and probability interpretations that rely upon these false solutions will quickly have to find new and fresh justifications. Or else they will die.

## The Finite World Argument

To ban everything that contains an unlimited number of entities from the realm of the possible, as this argument does, is both too drastic and too feeble at the same time. It is too drastic because if implemented universally all of mathematics as we know it would break down. Even the ancient Greek mathematicians knew that every finite entity could be expressed as an infinite sum of entities as well. For example, if I go from point $A$ to point $B$ this can be described as either a finite number of steps or as an infinite number of steps. A finite description would be to simply count the steps I need to take to go from $A$ to $B$. This adds up to the total distance between $A$ and $B$. An infinite description could be this: I first go half the distance to $B$. From there I go half the distance of what is left. From there half the distance from what is remaining, and so on. The entire walk is then described by the infinite series half the distance + one quarter of the distance + one eighth of the distance and so on ad infinitum. This infinite sum of course equals the full distance. The point is that whichever way my walk from $A$ to $B$ is described, the total sum must always be the same. Obviously, the real physical distance cannot be dependent on how it is described. However, if we obey the Finite World argument we need to say that some descriptions, the infinite ones, are unrealistic. For these a finite cap has to be imposed for the number of steps that actually can be performed in reality. No matter how big a number we choose as the finite cap, the capped sum of steps will never equal the full distance between A and B. So according to some descriptions I walk the full distance between $A$ and $B$, but according to others I cannot make the full distance. The Finite World argument thus makes the distance between $A$ and $B$ dependent on how it is described, which is absurd. It is easy to see that this example is not isolated but can be multiplied to every area of mathematics. If the Finite World argument is taken seriously, mathematics as we know it would break down. This is drastic.

The feeble thing is that the Finite World argument does not solve the problem it was set out to solve. We can easily construct a situation where no actual physical entities are infinitely many - and yet the St. Petersburg paradox is still present. Imagine a situation where we have a gambling hall with a number of different games. Some of the games are classical lotteries of various types. Others are variants of the St. Petersburg game with different payoff functions and fees for playing them. A set of players is invited to the gambling hall to try their luck at the different games and lotteries. The contestants are given an equal large amount of playing chips that can be used to pay the fees for the games in the hall. Each contestant is free to play as much or as little she wants at each lottery or game. If she decide to not play anything at all, that is fine too. If she wins any of the lotteries or games she will get her reward in 'winning chips' that can only be collected; they cannot be used for playing. The goal for each contestant is to have as many chips at the end of the day as possible. All playing chips that might remain, together with all the winning chips each contestant have won, are counted. The contestant who has the largest total collection of chips will get a nice prize - a car, say.

In this situation the Finite World argument is useless. The only physical prize present is a car, which is not infinite. The chips can be multiplied indefinitely and do not need to be physical, so no cap can be imposed on any of the payoff functions for any of the St. Petersburg games. This means that the expected value for each of the St. Petersburg games is exactly the same, i.e., infinite. But, obviously, it is better to play a St. Petersburg game with payoff function 1,000 chips, 2,000 chips, 4,000 chips, 8,000 chips, and so on instead of the original game with payoff function 1 chip, 2 chips, 4 chips, 8 chips, and so on, assuming the fee for the games are the same. However, the theory does not make any distinction between these games at all. This means that you are still stuck with the original St. Petersburg problem and the Finite World argument is of no use at all. This is feeble.

We have now showed that the Finite World argument can never lead to a real solution to the St. Petersburg problem. All solutions in this category are thus false.

## The Human Value Argument

The idea here is that ordinary humans do not think like mathematicians. In particular, ordinary men do not value large amounts of money in the way mathematicians do. For this reason ordinary humans do not have to obey the mathematical rules the theory produces. The theory of expected value is certainly correct mathematically, but it is simply not correct as a model for how ordinary men actually behave.

Ordinary men are driven by how happy they can get, not directly by how much money they can get. Sure, more money will probably make you more happy, but will twice the money you already have make you exactly twice as happy as you already are? Probably not. Twice as much money does not mean you will become twice as happy, and this psychological effect becomes even truer if your fortune is multiplied even more times. It seems to be some kind of law that the human appreciation of more money increases more slowly than what the actual amount of money possessed would indicate. A mathematical function that describes the decreasing additional happiness, or utility, you might get for increasing amounts of money is today called an utility curve.

If the utility, as described by a suitable utility curve, is inserted instead of the actual amounts of money you could win in the St. Petersburg game, the expected value, or rather the expected utility, becomes finite. This explains, according to this argument, why ordinary humans are not willing to pay what the mathematical theory demands in this case, but only a small finite amount.

This approach raises a number of new questions. Which mathematical function describes best the human utility of money? Should it be a universal curve for all men (except, perhaps, for mathematicians) or should it be different curves for different individuals? Can these curves be determined empirically? If so, how?

Interestingly, these questions have, over the years, not only been investigated thoroughly - the very idea of expected utility as a guide for human action has been hugely influential as a foundational concept. Almost all modern economic theory, decision theory and theory of rational choice rely on this concept. For example, in economic theory it has a prominent place in the "Law of diminishing marginal utility" and in decision theory in the "Expected utility hypothesis." The concept of expected utility is also a vital part of the Bayesian interpretation of probability.

Despite its fame, the concept of expected utility never solved the St. Petersburg problem that it was originally designed to solve. This is easily seen by the following example. In the original St. Petersburg game, instead of winning one dollar, two dollars, and so on, let the payoff function be in "utility units" instead of in money. Instead of winning two dollars you win the amount of money that exactly makes you twice as happy as you would be if you won one dollar, and so on. The expected utility will in this case be infinite, and the utility argument is of no use anymore as all payoffs already are utilities. Hardcore believers in expected utility theory, when confronted with this example, usually resort to some arbitrary Finite World argument to escape the embarrassment. However, the very reason utility curves were invented in the first place was to construct a solution to the St. Petersburg problem that did not have to resort to silly Finite World arguments.

Others, critical to the concept of expected utility, try to solve the St. Petersburg problem using a concept of risk. It is because we are afraid to risk too much of our money that we are reluctant to pay a lot for playing the St. Petersburg game, even if the theory tells us that we should. This is another Human Value argument. Only our imagination sets the limits on how many different Human Value arguments we can invent to solve the St. Petersburg problem. However, there is a simple way to show that all Human Value arguments must be wrong, even those not invented yet. If it was the case that the St. Petersburg problem could be solved by a theory of human valuation of money, be it expressed in utilities or risk or whatever, it would be impossible to give an account of the St. Petersburg problem when neither humans nor money are involved. But this is indeed possible.

Consider a membrane which during one second transmits one hydrogen molecule with probability one half, two hydrogen molecules with probability one quarter, four hydrogen molecules with probability one eighth, and so on. How much hydrogen gas can we expect to be transmitted through the membrane during one second?

This is exactly the original St. Petersburg problem but now without any reference to neither humans nor money. This means that all solutions that rely on how humans value money are indeed useless in solving the St. Petersburg problem. This simple thought experiment shows that the Human Value argument can never lead to a real solution of the St. Petersburg problem. All solutions in this category are thus false.

## The Reinterpretation Argument

Almost all aspects of the original theory have been reinterpreted just to avoid the St. Petersburg problem. The simplest one is merely a linguistic trick. Mathematicians do not talk about infinite expectations anymore. That way of talking is abandoned. What they say, instead, is simply that the expected value in that case does not exist. Something that does not exist cannot give you any advice, can it? Only expectations that do exist, i.e., is finite, can give you any advice. This trick resolves the St. Petersburg problem because the theory does not give you any advice at all as the expected value simply does not exist.

The problem with this escape route is that the distinction between when the theory will give you advice and not becomes arbitrary. This is because the difference between a situation where the expected value "exists" or "does not exist" can always be made arbitrary small. Consider for example the St. Petersburg game played with a real coin. It is an empirical fact that one of the sides of a real coin is always slightly more probable than the other side. The difference can be very small, but still there is always a difference. If we
happen to choose the side which is slightly less probable than $1 / 2$, as the side we repeat until the other side shows up and the game ends, we will end up with an expected value for the game that is finite and "exists." If we happen to choose the other side as the one that we repeat, we will end up with an expected value that "does not exist." Usually we do not know which side of a real coin that is the slightly more probable one-we just pick one of the sides at random as the one to repeat. So in half of the cases when we play this game the theory does have some advice to give us, while in the other half it does not. It is thus totally arbitrary when the theory can give us some advice and when it cannot.

To avoid this arbitrariness some reinterpret the theory even further so that the theory never gives any advice at all, whether or not the expected value exists. This does not lead to the same arbitrariness as before, but a theory that does not give any advice at all, in any case, is a little strange. Why do we have a theory in the first place if it does not have any practical applications?

Some other thinkers in this category say something really interesting. They claim that the expected value gets its entire meaning from the Law of Large Numbers, which is the name of the observation that the average gain will approach the expected value, or
average dice value against number of rolls

mean, after a large number of repetitions of a game. The only thing we mean when we say that a game is fair is that in the long run we will win as much as we will lose. Say that we play the game where we get one dollar for each dot that shows up when throwing an ordinary die. The expected value for this game is 3.5 dollars. The Law of Large Numbers will guarantee that we will approximately break even in the long run when playing this game over and over for a 3.5 dollars fee. In the limit when we play an infinite number of games we will break even exactly. See the graph above.

Expected values so defined do not say anything about single cases. Expected values for single cases are viewed as unreliable or even meaningless. This makes the very question what one should pay for the St. Petersburg game a meaningless question. Only if we play the game over and over is it possible to know what we should pay. And indeed, if we play the St. Petersburg game an infinite number of times we should actually expect to win an infinite amount of money, exactly as the expected value shows. Adopting this statistical viewpoint resolves the paradox, according to this view.

This idea is very clever. The Law of Large Numbers will actually guarantee that the expected value reinterpreted in this way will always keep what it promises. A gambler paying the expected value for any game will with probability one gain as much money as she loses in fees when the number of games goes to infinity. This resolution of the St. Petersburg problem does not lead to inconsistencies or arbitrariness. But it still has major drawbacks why it cannot be viewed as a correct solution to the problem.

It totally misses the original question, which is to answer what a single game is worth. In particular, we still have no clue what to pay for the St. Petersburg game if it is offered only once. This reinterpreted expected value can only answer a restated St. Petersburg problem: What is the fair fee for each round if we play an infinite number of St. Petersburg games?

In addition, we have no clue how many times we need to play any given game in order to have permission to use the expected value as a fair fee. If this problem is avoided by saying that one should always play infinitely many rounds to have permission to view the expected value as a fair fee, well then suddenly all games imaginable are dealing with infinities. The problems originally attached to games with infinite, sorry "not existing," expected values are now affecting all games, also those that never caused any problems before. This is hardly an improvement of the situation.

Even if this solution of the St. Petersburg problem is the best so far, it is still very far from a correct solution. As mentioned at the beginning, there is also one additional proposal in this category that we need to treat separately as it is not merely a reinterpretation but actually a bold idea by William Feller to change the very definition
of expected values (Feller 1945; 1950). Even for him the Law of Large Numbers is what ultimately gives expected values their meaning. He slightly restates the mathematical expression for the Law of Large Numbers. His new expression is a true generalization of the classical Law of Large Numbers because, whenever the expected value is finite, his generalized expression becomes the ordinary expression. But whenever the ordinary expected value is infinite, or "does not exist," his formula produces a series of variable fees. This idea is truly innovative. Never before did anyone call into question the implicit assumption that the fair fee for a game must be a constant.

According to Feller, the fair prize for the St. Petersburg game is

$$
\frac{1}{2} \log t
$$

where $t$ is the number of times the game is played and log is the logarithm to base two. See the graph below.


Let us say we decide to play the game 256 times. We should then be willing to pay 4 dollars per game according to Feller, as $\log 256$ is 8 . If we plan to play 4096 times the fee we should be willing to accept is 6 . By this we immediately see that when the number of games goes to infinity the fair prize also goes to infinity, which is exactly what the statistical reinterpretation of expected value says. But now, with Feller, we suddenly know a lot more what the game is worth even if we do not play infinitely many times.

This solution to the St. Petersburg problem is by far the best presented by anyone. But, unfortunately, it raises more questions than it answers.

First of all, is it his intention to replace the old concept of expected value with this new concept everywhere? It does not seem so. Instead, this concept seems to be tailormade for games of the St. Petersburg type, i.e., for which the ordinary expected value is infinite, or "does not exist." In that case his new concept falls into the same trap as the first reinterpretation idea we considered; it will be totally arbitrary when this concept and the old one should be used. The example with the St. Petersburg game using a real coin will apply equally well in this case.

Secondly, his concept is not applicable for one or even a few repetitions of a game. It is obviously not correct for a single game as $1 / 2 \log 1$ is 0 dollar, and we know by the construction of the game that it is at least worth one dollar, not nothing. This means that we know for sure that it is false for small $t$. He even admits this explicitly himself. But when, at what number $t$, is his formula beginning to be trustworthy? He does not say a thing about that. If we try to avoid this difficulty by simply say that $t$ needs to be infinitely large for the formula to start to be trustworthy, then we have gained nothing by using a new concept that allows for variable fees.

Thirdly, we might ask ourselves why it should be, in any sense, fair to pay any of the finite fees suggested by his formula. It is, after all, only in the limit when $t$ goes to infinity that his new concept obeys the Law of Large Numbers and the fees becomes "fair in the classical sense," as Feller puts it. For no finite value of $t$ is his concept fair in any sense. This means that even Feller falls into the same trap as the statistical reinterpretation; we have to play an infinite number of rounds in order to know that the suggested fee is a fair fee.

Despite his innovative and radical approach, we see that even Feller fails to solve the St. Petersburg problem. This concludes the task of showing that there still does not exist a single acceptable solution to the St. Petersburg problem. We can safely establish that the St. Petersburg problem is an open problem.

## Probability Interpretations

Initially, it was a good move to replace the central concept of the theory-expected value-with the concept of probability. Probabilities seemed to be totally unproblematic, which was not exactly the case with expected values, as we have seen. However, after a while, it became evident that the concept of probability is problematic as well. Classically, the probability of an event is defined as the number of favorable cases divided by the total number of cases. For this to work we need to find atomic cases which are equally likely. For example, the probability of getting an even number when throwing an ordinary die is three over six, because there are three favorable cases ( 2,4 or 6 ) and six equally
likely cases. But what does it really mean to be "equally likely"? It must mean that they are equally probable. But "probability" is the concept we try to define here and is of course forbidden to use before it is defined, or else the definition becomes circular.

To save the classical definition a new principle is introduced-the Principle of Insufficient Reason. It states that if you do not have sufficient reason to believe that one possible case is more likely than any other you are entitled to assign the same probability to each case. This formally removes the circularity in the definition of probability. However, this principle leads in itself to problems that are even worse-contradictions. Depending on the way we describe a situation, we will end up with different probability assignments for the same event. In addition, the principle cannot handle events where we have a continuum of outcomes. Even in these cases the principle leads to contradictions, as was first noted by Joseph Bertrand. We cannot have a principle at the core of our theory that leads to contradictions. It is evident that we need to give up this initial definition of probability entirely. It cannot be rescued. This is no good news at all. The concept of probability replaced the central concept of expectation just in order to give the theory a solid conceptual foundation. And now this solid core has evaporated into thin air, leaving a big conceptual hole at the heart of the theory. We are left with a mathematical theory that we have no clue what it is all about.

This situation needs to be solved in some way. But how? Very much as in the case with the St. Petersburg problem, a set of very different solutions to the problem emerges. In fact, it is because of the St. Petersburg problem people start to run in different directions when trying to find a solution to this problem. Depending on which strategy you settled for regarding the St. Petersburg problem, you will have different needs regarding the probability concept, and hence will end up advocating different definitions of probability.

If you believe in the Reinterpretation argument where expected values only are given a statistical interpretation as an average in a long run of games, you will end up with a frequentist interpretation of probability. According to this interpretation probabilities have no meaning for single cases, only a long-infinitely long-sequence of events from repeatedly perform an experiment can be attributed a probability. The probability is then defined as the relative frequency with which the outcome you want to measure occur in the sequence.

Notice how close this definition of probability is to the corresponding proposed solution of the St. Petersburg problem. In both cases it demands that we repeat the event we are interested in infinitely many times, no matter what event it is. Both concepts get their ultimate meaning from the Law of Large Numbers. Another interesting thing to note is that this definition turns the classical relationship between probability theory and
statistics upside down. Probability theory is now ultimately based on statistics and not the other way around.

If you are of a more radical type and go for the non-interpretation alternative where expected values have no moral implication at all, you will end up having the same view on the probability concept as well. Anything that fulfills the standard axioms of probability will be an admissible probability to you.

If you believe in one of the Human Value arguments you will end up being a Bayesian or adhere to a logical or epistemic interpretation of probability. In these cases the concept of probability does not derive its meaning from the Law of Large Numbers, but from another theorem in probability theory—Bayes' theorem. As it stands, Bayes' theorem is totally uncontroversial. It only becomes controversial when placed as the foundation for every application of the concept of probability, even in cases when the prerequisites of the theorem are not fulfilled. The idea is to try to rescue the Principle of Insufficient Reason in a way that does not lead to contradictions. This is accomplished via an ingenious idea. The subjectivity of the utility concept is here transferred to the probability concept itself. You are thus free to assign any probability you want to any event, and it does not have to be in accordance with anyone else's assessment of the same event. Instead of trying to create a theory that in itself is free from contradictions, the burden of consistency is placed on each individual. It is up to you to make sure that none of your probability assignments lead to contra-dictions! This is where the extended version of Bayes' theorem enters the stage. If you feed the theorem with your initial beliefs, what the theorem produces will always keep you safe from causing any internal inconsistencies. Your probabilities will of course still contradict other person's probabilities, but that does not matter. The only important thing is that your internal sets of beliefs are not contradictory.

Note how this definition puts the relationship between probability theory and its users upside down. Instead of giving advice to its users on how to play games, the users now, so to speak, give advice to the theory or feed the theory. The individuals hold the ultimate truth themselves and the theory of probability merely functions as a set of traffic rules for how the individuals should think so that they do not end up having two different thoughts that collide. What happens if you do not follow the traffic rules is obvious-you become irrational. And who wants to be irrational? This idea of absolutely rational agents is the basis for rational choice theory as well as for most of modern economic theory. In a free market the agents are assumed to act in an absolutely rational manner, i.e., according to exactly the same concept of rationality as used in Bayesianism.

While the Principle of Insufficient Reason only could handle situations where all cases were equally likely, Bayesian theory does not have this limitation. Your personal estimates
of probabilities for different events can have any values whatsoever-they do not have to be the same for each case. Mathematically, this is expressed with a prior probability distribution, as the Bayesians denotes a subjective valuation function. This function can have any shape you like; it does not have to be a constant function with the same value everywhere. For example, if you have a die in front of you that you suspect is loaded, you will naturally assign probabilities to the sides that are not $1 / 6$ for each side.

This brings us to another basic principle of the Bayesian philosophy. If you suspect that the die in front of you is loaded, you have to take this into account when you set up your prior probability distribution. If you do not, you will easily contradict yourself. In fact, you have to include exactly everything you know at every instant when you make assessments regarding probabilities. It is not that easy to be a Bayesian.

There are many variants of Bayesian probability, where some are called logical, epistemic, or objective Bayesian probability. However, the basic idea is the same but emphasis on what is important is placed at different places in respective philosophy. Objective Bayesians, for example, believe that there are objective ways to construct the prior distribution function, which removes the subjectivity from the theory. Logical interpretations stress the idea that probability, with the help of Bayes theorem, should be viewed as an extension of ordinary logic, namely the logic extended to uncertain statements. Epistemic probability stresses the importance of evidence as the basic force behind probability assignments.

What they all have in common is a desire to be able to assign probabilities to also single cases, something the frequency and statistical interpretations fail to do. Attempts to extend the frequency approach to account for single cases as well are usually called Propensity interpretations of probability. Physicists have always been happy with the statistical concept of probability for their needs—until Quantum Mechanics entered the scene. Now, suddenly, a physical theory made probability statements about single events, which the old statistical concept of probability never can give an account for. Karl Popper started the movement of creating a propensity probability concept. One central goal is to extend the statistical theory of probability so that it becomes compatible with how probabilities are used within Quantum Mechanics.

## Moving Forward

As we have seen, we have a strong historical and conceptual connection between the attempts to solve the St. Petersburg problem and all the current probability interpretations. We have also seen that the St. Petersburg problem is an open problem. But showing that
the St. Petersburg problem is still unsolved does not, by itself, prove that all probability interpretations are incorrect. It could be that one of the interpretations, while failing to solve the St. Petersburg problem, still has some justification as a probability interpretation. What we need is another example that explicitly shows how all current interpretations fail. Luckily, we have such an example. In recent decades another thought experiment has been discussed intensely, resulting in an even more confused debate than for the St. Petersburg problem. Also in this case, people are not able to agree upon what type of problem it is. Is it a problem within mathematics, logic, economics, psychology, or something else? Like the St. Petersburg problem it is a very easy problem to state, and yet no one has been able to solve it. However, the reason for this is very simple. None of the existing probability interpretations can be used to solve this problem, as we will see.

## The Two-Envelope Problem

Imagine that you have two sealed envelopes in front of you containing money. One contains twice as much money as the other. You are free to choose one of the envelopes and receive the money it contains. You pick one of them at random, by tossing a coin, but before you open it you are given the option to switch to the other envelope instead. Do you want to switch?

Obviously, the situation is symmetric, so it cannot make any difference whatsoever if you stick or switch. And yet, there is a clever argument that shows that you should switch. Assume that the envelope you picked contains A dollars. The other envelope must either contain $2 A$ or $A / 2$ dollars depending on if $A$ is the larger or smaller amount. You tossed a coin, so you know for sure that it is a 50/50 chance for either case.

There is a $50 \%$ chance that you get $2 A$ and a $50 \%$ chance you get $A / 2$ by switching. The expected value of switching is therefore $1 / 2 \times 2 A+1 / 2 \times A / 2$ which is $5 A / 4$, or 1.25A. This is more than $A$, what you already have, so you should indeed switch to the other envelope. But this clearly cannot be the case as the situation is symmetric. If you had selected the other envelope instead the same argument would have told you that you should have taken the envelope you now hold in your hand. The Two-Envelope problem is to spot and explain the flaw in this argument.

The frequentist or statistical concept of probability cannot even begin to solve this problem because it is evident from the setup that this is a single case. Propensity theories are of no help either. Those who think they can solve this problem are the Bayesians. The first Bayesian response is usually that the problem as stated is impossible to set up, as an infinite uniform prior distribution of money is implicitly required, and that is not
permissible according to Bayesian theory. Already this is a bit strange as many Bayesians have, in other contexts, argued that improper priors, as this is called, should indeed be allowed. To see why we end up with this improper prior, think about how the two envelopes could have been filled with money in the first place. If you pick an envelope that contains $A$ and it is equally likely that the other envelope contains $2 A$ or $A / 2$ it must mean that there were initially, before your were offered to pick an envelope, two envelope pairs $\{A / 2, A\}$ and $\{A, 2 A\}$ where both of them were equally likely to be picked as the pair you got in front of you. Only in this case is it equally probable that the other envelope contains 2 A as it is that it contains $A / 2$. But this must be the case for each of the possible amounts in all envelopes, in each possible envelope pair, so there must have been an infinity of envelope pairs for the person setting up the game to choose from: ... , $\{A / 8, A / 4\},\{A / 4, A / 2\},\{A / 2, A\},\{A, 2 A\},\{2 A, 4 A\},\{4 A, 8 A\}, \ldots$. Each of the pairs must have had an equal probability to be chosen as the pair you got in front of you. This produces the improper distribution we talked about, as every envelope pair will have probability zero and yet when summing them all they must add up to one.

However, this solution can be escaped by changing the Two-Envelope problem slightly, by introducing an explicit prior probability distribution for envelope pairs that is not an improper distribution, but a proper one. Such a distribution cannot be a uniform distribution why adjacent envelope pairs will have slightly different probabilities. However, when carrying out the calculations we will still get the conclusion that we should switch to the other envelope whatever we find in the first envelope. But as we know that this will always happen we do not have to open the first envelope we pick. We already know in advance that the other envelope is slightly better. This is truly paradoxical. And yet, the situation is as symmetric as before so any calculation leading to a difference between the envelopes must be false. In this case the Bayesians cannot blame us for implicitly assuming improper prior distributions anymore. The prior probability distribution here is explicit and proper.

To escape this trap the Bayesians usually still blame it all on the prior distribution. But not for being improper but for having an expected value that is infinite, or an expected value that does not exist if you will. This is a quite strange argument. The question if improper priors should be allowed or not has been discussed among Bayesians as long as Bayesianism has existed, but now suddenly it is not possible to use probability distributions within Bayesian theory which lack expectation. There exist whole families of standard probability distributions that are used every day by statisticians all over the world that lack expected value and variance. The Cauchy distributions, for example, is one such family of distributions. Bayesians are usually proud of being able to extend the
application of probability from the narrow set of applications the statistical or frequentist interpretation can offer. Here we see an example of the opposite trend. Important probability distributions that are totally unproblematic for frequentists to use are banned by Bayesians.

Most Bayesians are nevertheless happy with this explanation of the Two-Envelope problem. They are confident that new even more evil versions of the paradox will not appear. This is because a theorem shows that in order to produce the paradox the prior distributions need to have an infinite expectation. But a version of the paradox was already published early on by Raymond Smullyan where we cannot blame a prior probability distribution for being the culprit (Smullyan 1992). In fact, this version does not need any probabilities at all, so no probability distributions whatsoever, prior or otherwise, enters the scene.

Consider the same setup as in the original problem where you are given the option to pick one of two envelopes where one contains twice as much as the other. The following two plainly logical arguments lead to conflicting conclusions:

1. Let the amount in the envelope you chose be $A$. Then by switching, if you gain you gain $A$ but if you lose you lose $A / 2$. So the amount you might gain is strictly greater than the amount you might lose.
2. Let the amounts in the envelopes be $Y$ and 2Y. Now by switching, if you gain you gain $Y$ but if you lose you also lose $Y$. So the amount you might gain is equal to the amount you might lose.

The usual Bayesian response to this version is that this is another problem than the original because it does not include probabilities. This is a strange argument. If we can preserve the paradox while removing one concept that we initially thought was vital from the account, we have indeed learnt something. In this case, the Two-Envelope paradox is not dependent on the concept of probability, at least not a Bayesian concept of probability.

Incidentally, it is indeed possible to construct a probabilistic variant of the TwoEnvelope problem that neither include improper priors nor priors with infinite expectations.

## Jailhouse Clock

Imagine that you find yourself in death row in a prison in Texas for a crime you did not commit. You do not know when you are going to be executed. On the wall there is a
clock that you have noticed is a bit odd. It works properly except when the small hand is between noon and 1 PM, and between midnight and 1 AM. What should take one hour here always takes exactly two hours. Apart from this the clock works as normal. This means that the clock needs 13 hours to complete a full cycle.


A prison guard enters your cell and tells you that the time for your execution and another fellow prisoner has been determined. He hands over two letters with the execution orders containing the time of execution. You are free to pick any of the letters and the one you pick will determine when you will be executed. You pick one that says that you will be executed at 4 o'clock some day, but the actual day is not specified. Then suddenly the prison guard has a big grin over his face. He says that he is in a good mood today and want to give you an offer. If you want you are free to take the other execution order instead. It is a really good offer because the probability is exactly $1 / 2$ that your time left in life will be doubled, and with probability $1 / 2$ that it will be cut in half.

The guard is ignorant of what is stated in the letters. He only knows the procedure for how to pick times for execution from the clock on the wall. Each hour has an equal chance of being selected for an execution. Execution orders are always created in pairs where one time left to live for a prisoner is twice as long as the other. This is possible to do via the clock on the wall in a cyclical manner, due to its odd feature. On that clock, twice of 1 is 2 , twice of 2 is 4 , twice of 4 is 8 , twice of 8 is 3 , twice of 3 is 6 , twice of 6 is 12 , twice of 12 is 11 , twice of 11 is 9 , twice of 9 is 5 , twice of 5 is 10 , twice of 10 is 7 and twice of 7 is 1 . The pairs of letters presented to the prisoners are thus either $\{1,2\}$ or $\{2$, $4\}$ or $\{4,8\}$ or $\{8,3\}$ or $\{3,6\}$ or $\{6,12\}$ or $\{12,11\}$ or $\{11,9\}$ or $\{9,5\}$ or $\{5,10\}$ or $\{10$,
$7\}$ or $\{7,1\}$. Twelve different pairs are possible in total. The prior probability for each pair of letters to be selected is $1 / 12$. Because of this construction, when one letter in a pair is opened it is exactly as probable that the other envelope contains twice as much time left alive as it is half as much. In your case seeing the time stamp " 4 o'clock" reveals that the only possible pairs of letters are $\{2,4\}$ and $\{4,8\}$, and they are exactly equally likely to have been picked. The guard can therefore be absolutely certain that what he just told you is correct.

You are a Bayesian and know how probability theory can help you here. You have absolutely no reason to doubt that the other letter has either the time stamp " 2 o'clock" or " 8 o'clock" with probability $1 / 2$ each. By switching letter you will double your time left or cut it in half. What you can gain is thus twice of what you can lose, with an exact 50/50 chance for each outcome. Selecting the other letter will increase your chances of being released much more than it will decrease it. You know your lawyer needs every extra day she can get to reopen and win your case, before it is too late. So you would really need some extra time alive. On the other hand, of course, the situation is completely symmetric between the two letters you are given. No matter which one you would have opened first the other one would seem the more attractive. In particular, your fellow prisoner who got the other execution letter is also a Bayesian and reasons in the same way as you do. Both of you think it is a great opportunity to take the other letter so you end up swapping your execution letters.

None of the existing probability interpretations can solve this paradox.

## Two Old Problems

In the early 1650's, four gentlemen went on a three-day trip from Paris to Poitou, in the mid-west of France (Todhunter 1865). During the trip they discussed interesting philosophical topics such as the ontological nature of infinity, the existence of the infinitely small, the nature of the number zero and the existence of absolutely empty space. Above all, they discussed the nature of mathematics in general and its connection to reality. It was during these discussions that one of the gentlemen, Antoine Gombaud, alias chevalier de Méré, brought up two gambling problems that he thought strengthened his philosophical position. His stance was that mathematics is a very beautiful art on its own but cannot, in general, be trusted when applied to the real world. In particular when mathematical reasoning tries to embrace the mysterious concept of infinity.

To show that mathematics can lead to paradoxes even when no infinities are involved, he reveals a curious fact that he had discovered himself. Gamblers are interested
in calculating the so called critical number of a game. It is the number of times a game needs to be repeated in order to shift the odds from the gambling house to the player. For example, if you are offered to bet on a specified side of an ordinary die, you should only accept the bet if you are guaranteed to throw the die at least four times. Then you have more than a fifty-fifty chance to win the bet. The critical number for this game is thus four.

The strange thing Gombaud had discovered was that when playing this game with two dice the critical number is not what you would expect. As the number of possible outcomes increases by a factor six when playing with two dice instead of one, the critical number ought to be increased by a factor six as well. Six times four is twenty-four, but strangely enough the critical number when using two dice is not twenty-four but twentyfive. How could this be? According to Monsieur Gombaud this fact was nothing less than a scandal, as ordinary arithmetic apparently contradicts itself. One of the other gentlemen on the journey, Blaise Pascal, got upset by Monsieur Gombaud's view of mathematics as something beautiful but poorly connected to reality, and sometimes even contradicting itself. He decided to prove Monsieur Gombaud wrong as soon as he was back at home.

Monsieur Gombaud's problem is interesting. Ideally, the critical number really should increase by a factor six for every new die we add, for the reasons Gombaud devised. However, due to the discrete nature of a die the rule is not correct for the first few dice we add. As we add more and more dice the factor does indeed approach six. See the graph below.


For this particular game it is easy to calculate the critical number exactly for any number of dice. In general, however, this is not the case. Usually it quickly becomes an impossible task due to the vexing number of combinations of outcomes that need to be calculated and ordered. The ideal property of Monsieur Gombaud proves to be a handy tool for calculating the approximate critical number for games that need to be repeated many times to break even.

Unfortunately, Pascal never solved this problem, as he did not view it as interesting. In fact, he did not even understand the question. Instead he focused on the other gambling problem Monsieur Gombaud brought up during the trip. It was an old puzzle already then but new to Pascal, known as the Problem of points.

Two persons agree to put some money at stake and the goal of the game is to collect a specified number of points, say ten, alternatingly throwing a die. If they decide to quit playing for some reason after they have started to collect points, how should the stake be divided in a fair way? If they have collected an equal amount of points the stake is simply divided in half, but what to do if one of them is in the lead?

This problem had puzzled mathematicians and philosophers for over a century with no consensus on how to solve it. Pascal started to discuss this problem with his father's friend Pierre de Fermat via a series of letters. Initially Pascal was not sure at all that his solution to the problem was correct. However, when Pascal learnt that Fermat independently had arrived at the same division, albeit using other mathematical arguments, it made a huge impact on him. He quickly became convinced not only that their new principle was correct, but also that it could be applied to any type of decision problem. For instance, he devised a novel argument based on this principle for why one ought to believe in a god, today known as Pascal's Wager.

In modern terminology, their solution to the Problem of points amounts to the idea that each player should get a share of the stake that is proportional to their probability to win the game, had the game not ended. This share became known as the expectation or the expected value. Ever since its inception, this concept has been hugely influential in a number of different human inquiries.

## Fair Values

Why did Pascal and Fermat view the expected value as the fair division of the stake? Fermat apparently did not see the need for an independent justification at all. He seemed to think that it is mathematically obvious that his solution is the correct one. Pascal, however, tried to justify his solution by using a mathematical reasoning where each
player's fair share of the stake in the end can be reduced to a fair coin flip. To flip a fair coin to win your fair amount is seen as the quintessential fair game, and anything that can be reduced to this game ought to be fair as well. But to be reduced to a fair coin flip and be a fair coin flip are two different things. This is Pascal's mistake, which eventually led to the St. Petersburg problem and the messy philosophical situation we have today. However, no one at the time spotted this flaw in his argument. On the contrary, mathematicians all over Europe instead began to be interested in this new branch of mathematics, founded on the concept of expected value.

The same year the St. Petersburg problem was discovered the first version of the Law of Large Numbers was proved, giving the concept of expected value a much-needed theoretical support. It states that the expected value can be viewed as the average value for a game that is played an infinite number of times. Hence, if we play a game infinitely many times, we are mathematically justified to use the expected value as the fair prize. But if we do not happen to play a game infinitely many times, how can we justify to use the expected value?

After a large number of iterations of a game, the average value begins to fluctuate around the theoretical mean value, i.e., the expected value. Half of the time the average is above the mean and half of the time it is below the mean. This implies that if we play long enough, the probability is one half that we will end up as net winners and the other half that we will end up as net losers-if we pay the expected value each time we play the game. Just by iterating, any game will, in this sense, end up being equivalent to tossing a fair coin, which is the quintessential fair game. So, even if for a given game the expected value is not in any sense fair for a single or a few rounds, when repeated sufficiently many times, the series of rounds viewed as a whole will always be a quintessential fair game.

We see by this that even if the expected value from the outset is not a fair prize in any reasonable sense-just by repeating the game over and over we will arrive at a situation that models the quintessential fair game, i.e., tossing a fair coin once. That is, if the expected value is finite. If the expected value is infinite, as in the case of the St. Petersburg game, we will still come closer and closer to the "fair" value, which is infinity, but of course only from finite values. That is, only from 'below.' We will never reach a state where the accumulated gain will fluctuate around infinity, half of the time above infinity and half of the time below infinity. This is why the expected value is so strongly felt as being an unfair prize for the St. Petersburg game. It is not the infinite prize per se that is the problem, but the fact that the game cannot model the quintessential fair game no matter how many times we play the game. We can now define what a fair game is in the classical sense.

Definition A game is fair in the classical sense if and only if the probability is equally big for a net gain as for a net loss for all large number of rounds of the game.

Every game with a finite expected value can be turned into a fair game in the classical sense by assigning the expected value as the prize for the game. Note that no significance at all is put on how much we win or lose in the long run, only that the probability for a net profit is equal to the probability of a net loss. A net profit of one dollar is considered a 'win' as much as a net profit of millions of dollars. If we have a fair game in the classical sense we can guarantee that we will win or lose with equal probability in the long run, but we cannot, in general, say anything about how large the net gain or net loss will be. A closer analysis of the game at hand is needed for that. Hence, for the definition of fairness in the classical sense the size of the possible gain or loss is irrelevant.

However, we always need to perform a deeper analysis of the game at hand to know how many times we need to iterate the game in order to have reached the state when the game is fair in the classical sense. For some games the expected value is fair in the classical sense from start while for others we need to play an astronomical number of rounds. This is the meaning of the 'large numbers' in the Law of Large Numbers. But if this fifty-fifty chance for loss or gain in the end is the basic intuition behind the concept of fairness, why not use this upfront as the definition of a fair game?

Definition A game is fair if and only if the probability is equally big for a net gain as it is for a net loss.

From these definitions, we see that a game is fair in the classical sense only if a long sequence of games taken as a whole eventually becomes fair. But if a game is fair, it is automatically also fair in the classical sense. In the language of mathematical logic, 'fair' is a conservative extension of the old concept 'fair in the classical sense.' Hence, we have nothing to lose in adopting this new more general concept of fairness.

In mathematical language, an equal probability for gain or loss is called the median. The expected value is not a median but a probability weighted mean. What we noted above is that after a sufficiently large number of rounds played, the mean behaves exactly like a median-it actually becomes a median. This means that the median will approach the mean more and more the more the game at hand is repeated. It is in fact this medianlike property of the expected value in the long run that is the only valid justification for calling the expected value a fair prize.

## Solving the St. Petersburg Problem

If the approach to the median is the only reasonable reason to stick to the expected value - why not use the median upfront? The median is a fair prize by definition. For any game with a uniform distribution, the median and the mean will coincide. For a game with a non-uniform distribution, the mean and the median will not coincide in general. For example, the St. Petersburg game has an infinite mean (expected value) while the median (fair prize) is only 1.5 dollars.

If we play the St. Petersburg game twice, the worst-case scenario is that we win only one dollar each time, that is, two dollars in total. This will happen with probability $1 / 4$, because first we have to win one dollar with probability $1 / 2$ and then another with probability $1 / 2$, and $1 / 2 \times 1 / 2$ is $1 / 4$. If we have a little more luck we win one dollar the first round and two dollars the second round, in total three dollars, which has probability $1 / 8$, as $1 / 2 \times 1 / 4$ is $1 / 8$. Or, we win two dollars the first round and one dollar the next round, which also totals three dollars with probability $1 / 8$. Adding the probabilities for all these three worst case scenarios we get $1 / 4+1 / 8+1 / 8=1 / 2$. So, with probability $1 / 2$ we will win 3 dollars or less. Hence, with probability $1 / 2$ we will win 4 dollars or more. The median for playing the St. Petersburg game twice is thus 3.5 dollars. The fair prize per game is therefore in this case 1.75 dollars, which is slightly more than the fair prize for playing the game only once.

That the fair value increases is what we can expect as we know that the median must approach the mean, i.e., the expected value, the more rounds we play of the game. In this case we know that the expected value is infinite. The fair prize must therefore increase without bound for larger and larger sequences of repeated rounds of the game. In practice, this means that the more we are guaranteed to play the St. Petersburg game the more should we be willing to pay for the privilege to play the game. That a fair prize must, in general, vary depending on how many times we plan to play a specific game was realized already by William Feller, as we have seen.

In the graph on the next page we see how the fair prize per game increases as we are guaranteed to play longer and longer sequences of St. Petersburg games. As already mentioned, we know that this curve must increase without bound. But can we find an expression that approximates this curve? This is in fact possible using the same idea as Monsieur Gombaud used.

The more we play the game the more certain we are that we will win the most common prizes in proportion to how likely they are. If we play four times we can be somewhat sure that we will win the one dollar prize half of the time, that is in two of the four cases. This will give us two dollars. Ideally, we would expect to win two dollars in one

of the two remaining cases, four dollars in half a case, eight dollars in one fourth of a case and so on. However, this is clearly impossible, but the previous sentence describes exactly twice the St. Petersburg game played twice. The exact fair value for playing the standard game twice is 3.5 dollars as we already know. Therefore, the remaining two of the four cases contributes exactly $2 \times 3.5$ dollars. In total we have an approximate fair value of 2 dollars $+2 \times 3.5$ dollars for playing four times, or $1 / 2+1.75$ dollars per game.

If we play eight rounds, we will ideally win one dollar in four cases, two dollars in two cases and the last two cases is exactly like playing four times the St. Petersburg game twice. The approximate fair value is therefore 4 dollars $+2 \times 2$ dollars $+4 \times 3.5$ dollars, or $1+1.75$ dollars per game.

If we double once more and play sixteen times, we will arrive at an approximate fair prize of $3 / 2+1.75$ dollars per game. In general, the approximate fair prize per game when we play $t$ times is $1 / 2 \log (t / 2)+1.75$ dollars, where the logarithm is to base two. This can be rewritten as

$$
\frac{1}{2} \log t+1.25
$$

For large values of $t$, this expression gives an almost exact approximation of the true curve of fair prizes. In the graph on the next page this curve is shown in black. Thus, for large values of $t$ we can use the convenient formula above instead of calculating the exact fair value, which quickly becomes very complicated because of the vexing number of

cases to consider. Note how the 'ideal' black curve here plays the same role as the 'ideal' increasing factor of six in Monsieur Gombaud's gambling problem about the critical number.

If we solve the expression above for $t$, we get the relation

$$
\text { time }=\frac{4^{\text {money }}}{\sqrt{32}}
$$

where we have replaced $t$ with time to make the formula easier to remember. Whenever you are offered the opportunity to play the St. Petersburg game, try to remember this formula. Depending on the fee you have to pay to play the game you should not play unless you are guaranteed to, and have time to, play at least the number of times given by this formula. For example, if you are offered to play the game for a five dollars fee, you should not play unless you are guaranteed to play at least 182 times. If the fee is 20 dollars you will probably not have time to play the game even if you are given the opportunity to play the required amount of times.

This kind of advice sounds familiar. What is denoted 'time' in the formula above is exactly the same concept as the 'critical number' in Monsieur Gombaud's own gambling problem. We can thus conclude that this idea is far from new. In fact, it comes natural to most people. In clinical studies where people have been asked what they are willing to put at stake for playing different games, the concept which best fits the empirical data is the fair prize, that is, the median (Hayden and Platt 2009).

## Solving the Two-Envelope Problem

Unfortunately, the solution of the St. Petersburg problem does not solve the TwoEnvelope problem. If we replace all expected values by fair values, the latter problem remains, as is seen in the Jailhouse Clock scenario. Incidentally, this proves that these two problems are totally unrelated. The opposite stance, that they are closely related or even just variants of the same problem, is quite common among the commentators to the Two-Envelope problem. According to this view, the true solution of the St. Petersburg problem would automatically resolve the Two-Envelope problem. Now we see that this is not the case.

The Two-Envelope problem is superficially similar to the paradoxes invented by Joseph Bertrand in the nineteenth century that helped to kill the classical interpretation of probability. However, the Two-Envelope problem goes deeper as it also shows that the Bayesian interpretations are wrong. According to Bayesian philosophy 'uncertainty' or 'lack of knowledge' can always be modeled by a probability distribution in a way that does not lead to contradictions. But we can construct an explicit probability distribution describing a certain state of 'uncertainty' for the Two-Envelope problem that does indeed lead to a contradiction (Broome 1995). This shows that the Bayesian philosophy is false.

The underlying world view motivating the Bayesian philosophy is determinism. If the world is deterministic only our 'lack of knowledge' can be the reason for why we are uncertain about what will happen. It is always, however, possible to learn more about the situation at hand and reduce our uncertainty. For example, if we flip a coin and we have no clue which side will come up, we say that chances are fifty-fifty for either side. But if we know more about the actual coin or learn the physical details on how it is flipped, it is always possible to make a better guess. If we have total knowledge of the physical situation at hand we can predict with certainty which side will come up. Hence, for determinists it is natural to equate probability with lack of knowledge. Total lack of knowledge leads to a uniform probability distribution among the possible outcomes. Partial knowledge leads to some other probability distribution, which completely describes the state of knowledge. If we have total knowledge everything is certain and we have no need for probabilities. Or equivalently, all probabilities are either zero or one. We know for certain if something will happen or not.

To use probabilities as an irreducible entity in a fundamental physical theory is thus unthinkable for a determinist. If the theory really is fundamental it cannot rely upon a concept that is synonymous to 'lack of knowledge.' That is just another way to say that the fundamental theory really is not fundamental at all. There must exist some even more fundamental theory that explains the apparent randomness. This was exactly the
view Albert Einstein held regarding the new physical theory describing the very small, Quantum Mechanics. It is held to be a fundamental theory that nevertheless relies upon irreducible probabilities.

As a determinist, it was evident for Einstein that Quantum Mechanics cannot be a fundamental theory. To convince even non-determinists he developed, together with two coworkers, a philosophical argument that is now known as the EPR argument (Einstein et al. 1935). The argument uses two spatially separated particles that are connected in a special "spooky" way that Quantum Mechanics permits. Einstein viewed the connection as spooky because the particles seem to keep track of each other through space and time in an inexplicable way. For example, if a property like spin is measured in a particular direction for one of the particles the other particle always has a spin in the opposite direction. This would not be strange at all if their spin directions were predetermined and simply set in different directions from the beginning. But that is not how it is according to the theory. According to Quantum Mechanics, the outcome of the first measurement is completely random and not predetermined at all. How the other particle can know the outcome of a completely random event far away and adapt its properties accordingly is what is called "spooky action at a distance" in the EPR paper.

As soon as one of the particles has been measured, we know for certain the outcome of the corresponding measurement of the other particle. A property that can be predicted with certainty must imply that the property is real and exists objectively, even before we measure it. But the theory explicitly denies that this property could be real and existing before the measurement. According to Einstein, this clearly shows the incompleteness of Quantum Mechanics. A complete theory has something corresponding to every element of reality. Any property in the natural world that can be predicted with certainty, i.e., with probability equal to one, is an element of reality. This is how 'an element of reality' is defined in the EPR paper. As Quantum Mechanics cannot account for some quantum properties, which clearly are elements of reality according to this definition, the theory must be incomplete. This is the EPR argument.

The EPR argument is flawed because the reasoning is circular. If the world is deterministic we already know that Quantum Mechanics is incomplete, so the EPR argument cannot assume that. But the interpretation of probability used in the definition of "an element of reality" is Bayesian—probability as a measure of how complete our knowledge of a situation is. And, as we know, this interpretation only makes sense in a deterministic world where every particle has well determined properties all the time. So instead of giving a definite proof of the incompleteness of Quantum Mechanics, as was the intention, the EPR argument only shows that Quantum Mechanics is incomplete if
we assume that Quantum Mechanics is incomplete.
In fact, what the EPR thought experiment really shows is quite the contrary. It shows that irreducible true randomness indeed exists. There is a well-known theorem, called the no-communication theorem, which says that the coupled particles in the EPR setup cannot be used to send information faster than light. The reason we cannot utilize the "spooky action at a distance" for actual communication is because the first measurement we perform is totally random. If it were not completely random we would actually be able to send information faster than light. Indeed instantaneously. Not only is superluminal communication forbidden in the theory of relativity, but the very concept of 'instantaneity' is totally alien to it. So, in order to comply with Einstein's own theory of relativity, the EPR experiment shows that the world actually is indeterministic and not deterministic, which Einstein, Podolsky, and Rosen implicitly had assumed.

Knowing that probability must have to do with randomness and that our physical world is in fact indeterministic, does this by itself solve the Two-Envelope problem? It is an interesting fact that it does not. In the Jailhouse Clock scenario, for example, we can specify explicitly how to randomly select the pair of execution orders to be presented to you. The guards can simply use a twelve-sided die as the random generator and select one of the twelve possible pairs depending on what the die shows. Then you flip a coin to decide which of the two execution orders to choose. The Jailhouse Clock problem, however, prevails. This is why the Two-Envelope problem is deeper than the paradoxes presented by Joseph Bertrand. They all disappear when a random procedure like this is specified explicitly.

Which additional idea is needed to solve the problem? We know that the world is indeterministic because there are genuinely random events. But what does "random" really mean? In a deterministic world when repeating an experiment the outcome must always be the same. If it is not the same, we have not repeated exactly the same experiment. This is not the case in an indeterministic world. Here we can repeat exactly the same experiment and the outcome can still be different from one performance to the next. This is what we mean by fundamentally random events. So, every experiment in an indeterministic world leads to a set of random events. But is the opposite implication true? Does a set of random events imply an experiment?

Let us investigate this question in the context of the Jailhouse Clock scenario. You look at the letter with the execution order you have picked that says that you will be executed at 4 o'clock some day. Then, suddenly, you are given the opportunity by the prison guard to choose the other letter instead. This event is a random event where we do not know what the experiment is. Either the time stamp " 4 o'clock" is essential to
why you got the opportunity to switch letter by the prison guard, or it is not. As this situation will only happen once in your life you have no way of knowing what would have happened if some other time stamp would have shown up in the letter you first picked.

There are two options. Either you are given the opportunity to switch to the other letter whatever timestamp you got first, or you are not. In other words, either the timestamp " 4 o'clock" is part of the description of the experiment, or it is not. If it is part of the description of the experiment you are only given the opportunity to switch in case you found the execution order with timestamp " 4 o'clock." In this case it is advantageous to switch as the fair value of your time left alive is increased using the other execution order. If " 4 o'clock" is not part of the description of the experiment it must instead be one of the twelve possible outcomes of the experiment and you are given the opportunity to switch whatever time is stated in the first letter. If this is the case, there is nothing to gain by switching to the other execution order.

By this we see that a set of random events by itself does not imply a unique experiment. Sometimes a set of events is compatible with more than one experiment. Moreover, this also shows that attaching probabilities to events when the experiment is not defined leads to contradictions. 'Experiment' is thus a more fundamental concept than randomness.

Definition An experiment is a complete set of instructions on how to properly repeat something.

Note that this definition does not make any distinction between real experiments and thought experiments. It does not matter if the experiment is actually performed or not. It can occur many times, only once, or never. The set of instructions is always the same. If no experiment is defined, it is impossible to talk about probabilities. It leads to contradictions as we have seen. The experiment is thus the important concept here and not the random event itself. Random events cannot be attributed probabilities without a reference to an experiment. This is the key observation to understand the concept of probability.

Definition A probability is a measure of the relative occurrence of an outcome of an experiment, would the experiment be repeated indefinitely.

Note how this concept does not make any distinction between theoretical probabilities and actual measurable probabilities. This is because it relies upon the
concept of experiment above which does not distinguish between actual experiments and thought experiments.

This concept solves all probabilistic versions of the Two-Envelope problem. Obviously, Smullyan's non-probabilistic version of the Two-Envelope problem cannot be solved by a new concept of probability, as no probabilities enter that problem. Smullyan's paradox instead shows something very interesting about the logical structure of the world.

We cannot use our probability concept to solve Smullyan's problem but we can use the concept which is the logical basis for it, our concept of experiment. The two conclusions derived in the problem, that the amount we might gain is greater than what we might lose, and that the amount we might gain is equal to what we might lose, respectively, are only possible because different experiments are used to derive the conclusions. In the first case, the amount $A$ is part of the experiment while in the second case neither $Y$ nor 2 Y are part of the experiment. In the first case, the outcomes of the experiment are 2A and $A / 2$ while in the second case the outcomes are $Y$ and 2 Y . This is the reason we get different conclusions. Our concept of experiment thus solves Smullyan's non-probabilistic version of the Two-Envelope problem.

## Consequences for the Measurement Problem

We have shown that in order to avoid contradictions we always need to explicitly specify what our experiment is and clearly distinguish that from the results we get from performing the experiment. Not only when probabilities and random events are involved, but always.

This has interesting consequences for understanding the so-called Measurement problem in Quantum Mechanics. The Measurement problem is how to explain why measurements play such a peculiar role in Quantum Mechanics. As long as we do not measure a quantum system, the system evolves in a deterministic fashion. But as soon as we do a measurement the deterministic progression collapses and we get a random result that is in accordance with what the theory predicts. The problem is that this "collapse" of the deterministic evolution is not explained or motivated in the theory at all. Moreover, no one knows for sure what a "measurement" really is. No definition whatsoever is given by the theory. It is a mystery how this concept, which is in no way a part of the theory, can nevertheless be totally crucial for the application and therefore success of the theory. To explain this is called the Measurement problem.

If we try to understand Quantum Mechanics using a deterministic ontology, we will fail, as we have seen. In a deterministic world every object has exact positions
and properties all the time. But the world is indeterministic as we have shown, so we need to adopt an indeterministic ontology instead. In an indeterministic world, there are genuinely random events that cannot be reduced to more basic underlying random events. But the only way a genuinely random event can exist is because it is the result of an experiment. If we do not have an experiment, we cannot have a random event. This is true on a basic logical level of reality for which the physical reality has to comply. Nature has solved this problem by letting quantum systems evolve in a deterministic but unreal manner when no experiment is defined. Or rather, between the start of an experiment and the end of the experiment. As soon as the experiment ends, which is usually called a measurement, a random event will occur that is compatible with the experiment.

If an experiment is changed, the possible random outcomes are changed as well. To investigate the state of a particle in the middle of a genuinely random experiment is impossible, simply for the reason that such a measurement changes the original experiment. An experiment cannot have two different starting points at the same time. This follows from the definition of an experiment. As soon as the particle is interrupted, it marks the end of an experiment and a new experiment begins where the previous starting point is completely forgotten.

As the change of an experiment happens at a logical level of reality, the change has to be instant. This explains the instant change in the EPR setup. Nature's detection of which experiment is at hand is always instant. If it was not instant, it would be possible to arrange a situation where two experiments were defined at the same time for the same set of particles, which would easily lead to a logical contradiction.

The new probability concept, which we can call the experimental interpretation of probability, also explains why a "certain" prediction of an outcome, as in the EPR argument, still does not mean that the particle is in that state prior to the measurement. A measurement is a physical interaction which either has occurred or has not occurred. If it has not occurred, the particle is still in its unreal but deterministic state. As soon as the measurement happens the particle's property that is measured becomes real. The fact that we can with certainty predict the outcome of an experiment does not change this fact.

## Conclusions

At first, it might seem like a good thing that we allow for many different interpretations of the probability concept. But if that is the case, it must logically be a good thing to have many different interpretations of any scientific concept like force, distance, time, electric
current and so on. It is easy to see that if we had that, science would hardly be possible. There is no legitimate reason to claim that probability is special, in any sense, from any other scientific concept. Either it is good for all concepts to have many contradicting interpretations or it is good for none. It is clearly the latter case.

Allowing for only one interpretation of probability will have a major impact on both natural science and the human sciences. There will be no reason to use totally different types of mathematics in economics and physics, for example, as is currently the case. As this paper shows, both can instead start to use the same concept of probability. If adopted, this will have a great unifying effect across different disciplines.

## References

Broome, John. 1995. "The Two-Envelope Paradox." Analysis 55 (1): 6-11.
Dutka, Jacques. 1988. "On the St. Petersburg paradox." Archive for History of Exact Sciences 39 (1): 13-39.
Einstein, Albert, Boris Podolsky, and Nathan Rosen. 1935. "Can quantum-mechanical description of physical reality be considered complete?" Physical Review 47 (10): 777-780.
Feller, William. 1945. "Note on the Law of Large Numbers and "Fair" Games." The Annals of Mathematical Statistics 16 (3): 301-304.
__ . 1950. An Introduction to Probability Theory and Its Applications, Volume I. New York: Wiley.
Hayden, Benjamin and Michael Platt. 2009."The mean, the median, and the St. Petersburg Paradox." Judgment and Decision Making 4 (4): 256-272.
Jorland, Gérard. 1987. "The Saint Petersburg Paradox 1713-1937." The Probabilistic Revolution, Volume 1: Ideas in History, ed. by L Kruger, LJ Daston, and M Heidelberger, 157-190. Cambridge: MIT Press.
Samuelson, Paul. 1977. "St. Petersburg Paradoxes: Defanged, Dissected, and Historically Described." Journal of Economic Literature 15 (1): 24-55.
Smullyan, Raymond. 1992. Satan, Cantor, and infinity, and other mind-boggling puzzles. New York: Alfred A Knopf.
Todhunter, Isaac. 1865. A history of the mathematical theory of probability from the time of Pascal to that of Lagrange. Cambridge and London: Macmillan and Co.

