

Received July 25, 2020, accepted August 5, 2020, date of publication September 1, 2020, date of current version September 14, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.3020928

A Framework for Probabilistic Decision-Making Using Change-of-Probability Measures

SAEDE ENAYATI¹, AND HOSSEIN PISHRO-NIK¹, (Member, IEEE)

Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, MA 01002, USA

Corresponding author: Hossein Pishro-Nik (pishro@umass.edu)

This work was supported in part by NSF under Grant CNS 1932326.

ABSTRACT Probabilistic decision-making is a fundamental problem considered in many disciplines from engineering to social sciences. In this article, we address decision-making in contexts where the law of large numbers (LLN) does not apply. Non-LLN regimes include almost all high-impact decisions. The rise of artificial intelligence (AI) decision making is further increasing the importance of developing principled approaches for such problems. In this regard, we first introduce a method called bounded expectation (BE) to apply the accepted principle of ignoring negligible probabilities. We show that BE provides some satisfactory results and insights into some decision-making problems. Pointing out some shortcomings of BE, we then turn to a much more general setting, using change-of-probability measures. We show that the proposed approach can be considered a generalization of expected utility theory (EUT) from two different perspectives. First, the approach converges to EUT as the number of repetitions grows. Additionally, when the fundamental distortion parameter, ϵ , is set to zero, the proposed theory reduces to EUT. We then propose a systematic approach to applying the developed framework to non-LLN decisions. Finally, through a real-world example, we compare the decisions made with the proposed method and the conventional methods. It is speculated that due to the complexity and multidimensionality of decision-making under non-LLN regimes, the presented ideas can potentially lead to considerable further research, some of which is discussed in this article.

INDEX TERMS Decision-making, decision theory, expected utility, probability, risk analysis, St. Petersburg paradox, statistics, utility theory, uncertainty.

I. INTRODUCTION

Probabilistic decision-making has been studied in different disciplines, such as engineering, computer science, philosophy, economics, and other social sciences. In a typical scenario, an intelligent agent (human or AI), analyzes its environment and takes actions to achieve some goals. When the probabilities of the potential outcomes are known or could be estimated, such decisions are referred to as decisions under risk [1]. While such decisions are already central in decision theory, their importance is increasing as the role of AI becomes more prominent. The output of machine learning (ML) algorithms can usually be interpreted in terms of probabilities. For example, a classification algorithm normally outputs the probability that the input belongs to a certain class. Therefore, many AI-based decisions are probability-based decisions, i.e., decisions under risk.

The associate editor coordinating the review of this manuscript and approving it for publication was Shyi-Ming Chen.

Another major shift is that an increasing number of *high-impact decisions* are being made by AIs. This makes the issue of having a sound and rigorous foundation for probabilistic decision-making even more important than before. There has been significant and influential work in decision theory and related fields. A very common approach is the principle of maximizing expected utility. In this context, an agent makes decisions based on a utility function that is defined carefully to best represent the real “value” of potential outcomes. For example, in the context of reinforcement learning, typically the agent’s goal is to find policies to maximize the expected reward [2]–[9].

The principle of maximizing expected utility, while very useful in many contexts, has known limitations [10]–[12]. Somewhat implied in the axioms of the theory lies the assumption that decisions are made in environments where the law of large numbers (LLN) holds in some way. In fact, the very notions of probability and expected value are essentially asymptotic. Probabilities and expected values are empirically what is observed in the long run. Thus,

a very important limitation arises when an agent is making a *high-impact decision that cannot be aggregated into a large set of similar decisions*. In such cases, we cannot invoke the LLN to justify application of the maximum expected utility principle.

This article takes a fresh look at probabilistic decision-making with a focus on high-impact decisions for which the LLN cannot be used. We provide a general framework based on *change-of-probability measures* for making probabilistic decisions which can be applied in both LLN and non-LLN regimes. We will provide evidence that the proposed approach can provide consistent and satisfactory results in some practical settings. The proposed theory can be considered a generalization of expected utility theory (EUT) along two different dimensions. First, we show that as an agent repeats an action, the proposed theory converges to the expected utility theory. Nevertheless, the theory can produce results for any number of repetitions of an action. Therefore, the proposed theory can be considered a *nonasymptotic* theory for probabilistic decision-making.

There is a second aspect in terms of which the proposed framework can be considered a generalization of EUT. A central parameter in the proposed framework is the distortion parameter ϵ . Setting ϵ to zero reduces the proposed approach to EUT. We then propose a systematic approach for decision-making under non-LLN regimes by considering how preferences change as the value of ϵ changes.

A key insight in the proposed approach is that when an agent faces nonrepeatable decisions, it might be beneficial to place more weight on the most likely outcomes. Nevertheless, it is not clear how to do this in a *principled way*. This will be a main focus of this work. It is emphasized that the problem of decision-making under non-LLN regimes is a multifaceted problem, with each decision having unique aspects. Hence, it is most likely that no single approach can provide satisfactory results to all problems. Therefore, the presented work here could potentially lead to considerable further investigation. For example, combining the proposed framework with risk management techniques that mostly focus on the tails of the distributions (e.g., [13], [14]) might be a promising approach.

While the proposed theory can potentially be used in any probabilistic decision-making context, its importance could potentially be magnified by the prevalence of AI in the future for three reasons. First, as a greater number of high-stakes decisions are made by AI, it is crucial to have a formal and rigorous basis for such decisions. Second, as the output of many ML algorithms can be used to estimate probabilities, decisions under known probabilities are becoming more prevalent. Finally, the proposed theory is normative (as opposed to descriptive). Again, this is more suitable for AI applications as machines can be programmed to make optimal probabilistic decisions without being impacted by psychological biases and flaws.

A. RELATED WORKS

Probabilistic decision-making has been studied in many different disciplines. The rapid growth of AI and its applications has led to high-impact AI-based decision-making becoming more prevalent in the systems and technologies of different fields, such as medicine, autonomous vehicles, network and national security, safety and privacy, and business, [15]–[25].

As mentioned before, expected utility has been extensively used as the decision criterion. Furthermore, there are many works expanding and improving upon expected utility theory. A large collection of works are descriptive and focus on how humans make decisions [26]–[28].

There are also many normative works, such as [29]–[33]. Some works, such as those elaborating the risk-weighted expected utility approach, look at the risk aversion or risk-seeking attitudes of agents [33]. Other works focus on theories of how utilities are compared, such as rank-dependent expected utility theory [34], relative expectation utility theory [35], and cumulative utility theory [36]. As will be clear later on, it is possible to combine such methods with the proposed technique in this article as they address different aspects of decision-making.

All the abovementioned references represent significant progress in this broad field and indicate the importance of probabilistic decision-making. This article takes a fresh look at this problem from a different angle: It aims to develop a theoretical framework based on change-of-probability measures that addresses the fundamental observation that not all decisions can be aggregated into a large set of similar decisions (non-LLN regime). This is a crucial point, as many high-impact real-world decisions and incidents can be put in this category [37]–[44].

Change-of-probability measures have been applied in other contexts, such as finance [45] where the concept of a risk-neutral probability measure is used for the purpose of option pricing; communication [46]; signal processing [47]; and fuzzy measure theory [48]. Nevertheless, these applications are within a different framework, and the goals and mathematical constructions in them are different from the ones we consider here.

Historically, the St. Petersburg gamble appears to have been the first problem where mathematical expectations failed to evaluate the game's value in a rationally acceptable way. Hence, for more than 300 years, many efforts have been made to address the paradox [49]–[56].

B. CONTRIBUTIONS AND ORGANIZATION

Our contributions can be summarized as follows:

- This article provides a theoretical framework for probabilistic decision-making problem that can be applied to both LLN and non-LLN regimes. The proposed change-of-measure is determined, in a *principled way*, as a function of the involved random variables. Such a construction allows us to ensure that the decision-making policy satisfies important

requirements. To the best of our knowledge, this is the first paper that introduces such an approach.

- We introduce bounded expectation as a special case of the proposed method and show that the method can provide satisfactory solutions to some problems such as St. Petersburg paradox.
- Based on the proposed construction, we prove its desirable properties for example, the fact that the proposed method converges to the expected utility method if the actions are repeatable. The main value of the proposed approach, however, lies in its usefulness for situations in which actions are not repeatable or are repeatable only a few times.
- A key contribution is that a systematic approach for applying the proposed method is provided. The proposed method results in satisfactory answers and insights in the investigated examples.
- Finally, by turning our attention to a very popular real-world problem, i.e., angel and venture capital investment, we show that our proposed method provides a reasonable solution that, due to the benefits of being systematic, can be programmed to allow AI agents to make these kinds of decisions.

The paper is organized as follows. In Section II, we provide some discussions to better explain the motivation of the work. Next, before providing the general theory, we first present a special case of our general framework called bounded expectation in Section III. The reason for this choice is to use a simple and concrete example to focus on some important insights without becoming bogged down in mathematical details. We then formally present the theoretical framework in Section IV. There, we formally prove the ideas and concepts discussed in Section III in a much more general setting. Examples of more sophisticated and powerful change-of-measures as well as a systematic approach to decision-making under non-LLN regimes are provided in Section V.

II. MOTIVATION

To better motivate the discussion, we consider two examples:

Example 1 (St. Petersburg Paradox): Consider the well-known St. Petersburg paradox. An agent is offered the following gamble: a fair coin is tossed repeatedly until a head is observed for the first time. If k is the total number of resulting coin tosses, the agent is received $X = 2^k$ units of utility. The question is how much (in utility units) is this gamble worth to the agent?

Since the agent receives 2^k units of utility with probability $\frac{1}{2^k}$, the expected value is infinity:

$$\begin{aligned} E[X] &= 2 \times \frac{1}{2} + 2^2 \times \frac{1}{2^2} + \dots \\ &= \infty. \end{aligned}$$

This result is clearly controversial as has been observed by many mathematicians [49], [57]. For example, the probability that the agent wins more than 32 dollars (for simplicity, let

us replace utility units by dollars from now on) is only 3%, yet the calculation suggests that the gamble is worth infinity. If the agent is considering whether to pay a large sum for this game, then this game will be a high-impact decision.

To better understand the crux of the problem, consider a slightly different scenario, where the agent can choose to repeat the game *as many times as she wants* and for each play she pays c dollars. Let \bar{X}_n be the average amount received by the agent after n repetitions; then, we have for any $c \in \mathbb{R}^1$

$$P(\bar{X}_n \geq c) \rightarrow 1, \quad \text{as } n \rightarrow \infty.$$

In other words, no matter how much the agent pays for each game, the agent *eventually* wins more than what she pays (assuming c is kept constant). Therefore, in this version of the game, it is not irrational for the agent to pay a large fee to play each game. Thus, the amount that the agent should be willing to pay for each game should depend on the total number of times the agent is allowed to repeat the game.

The above issue is not limited to infinite-mean random variables. In a typical decision-making scenario, an agent might be faced with a one-time high-impact decision where with a very small probability (say $\frac{1}{1000}$), the agent will receive a very high reward but otherwise will receive a negligible or a negative reward. In such cases, the expected value of the reward might be very high, while it seems very unreasonable to place a very high value on such gambles since the agent is almost sure she will not receive the large reward.

Example 2 (Court Dilemma): As a second example, let us consider a plaintiff in a legal case who is offered a settlement in which she will receive one million dollars. Her lawyer estimates that she has a 35% chance of winning the case, whereupon she would receive ten million dollars, but she will receive nothing if she loses. Let us compute the expected utility for each option. Let us assume one million dollars has 10 units of utility, while ten million dollars has 40 units of utility (consistent with the diminishing effect of marginal utility [60]). Additionally, if the plaintiff goes to court and loses, the resulting outcome is not zero, as there is a psychological effect in terms of disappointment and regret. Let us thus assume that this outcome yields a utility of -5 .² Let X be associated with accepting the settlement and Y with going to the court. The expected utilities are

$$E[X] = 10,$$

$$E[Y] = 0.35 \times 40 - 0.65 \times 5 = 10.75.$$

We observe that going to the court actually yields higher utility! Nevertheless, this does not seem to be a wise choice:

¹This can be concluded from the version of the LLN extended to infinite-mean random variables; see for example, [58], [59].

²Two points: First, there is also some element of regret in accepting the settlement as the plaintiff may wonder whether she could have won the larger prize, so the 10 units of utility in that case is assumed to be computed with this effect taken into consideration. Second, the exact values of utilities here are not very crucial: it is clear that different people assign different utilities to outcomes. Nevertheless, the phenomenon being discussed can often be observed. In other words, you can change the monetary rewards so that the resulting utilities are given by the values assumed in this example.

with 65% probability, it will result in the worst possible outcome.

Again, here we observe that this case could be another example of a high-impact event that cannot be simply aggregated with other decisions the plaintiff makes in her life, so the LLN cannot be used to justify the expected utility approach. Of course, if the plaintiff were extremely wealthy, the story would be a different. In that case, this decision could be simply aggregated into her other financial decisions, and in that case, the LLN could be used to justify the expected utility approach.

The crucial point is the following: In real-life decision-making, there are scenarios where decisions cannot simply be aggregated into a large set of similar decisions. This is, for example, the case when a very high-impact decision is being made where the outcome might have large consequences. The issue is that such decisions are not governed by LLN, as they are not repeatable (or are only repeatable a few times). For such cases, the very notions of probability and expectation have limited use, as they are inherently asymptotic.

The fundamental question then becomes the following: Can we provide a theory that addresses the above issue? Such a theory might evaluate a decision differently based on how many times an agent is going to face similar decisions in total, e.g., how many times the action can be repeated. More specifically, as the number of repetitions grows, the values converge to the expected value, but the crucial value is in the finite repetitions.

This article aims to answer the above question. As we will see, the proposed theory based on change-of-probability measure provides a promising framework to achieve the above goal. In the cases we considered, the method produces results that are consistent with the decisions typically expected by decision theorists. The method can also explain several phenomena that have been empirically observed. For example, applying the theory to right-tailed distributions, we can explain the underlying dynamics behind the venture capital industry. Applying the theory to left-tailed distributions, we can gain insight into the robustness and fragility of decision policies. Finally, the theory provides satisfactory results for the two examples provided above.

Obviously, the general question of making high-impact decisions has many different aspects, and most likely, a single framework cannot provide all the answers. More likely, a combination of different approaches could provide the most satisfactory result. In that regard, the proposed framework can be considered a step toward achieving such an important goal. One positive aspect of the proposed method is that it can easily be combined with other methods such as those developed for risk management.

Finally, it is worth noting that in evaluating an action, there are two important parts: probabilities and utilities (values of rewards). In this article, we focus on probabilities. It is assumed that the agent can appropriately assign utilities (rewards).

III. BOUNDED EXPECTATION

In this section, we introduce a specific change-of-probability measure called bounded expectation (BE) and discuss how it can be used in decision-making. BE is a very simple version of the proposed theory and is not perfect. It does not fully enjoy the potential advantages of metrics that can be built using the proposed change-of-probability measure approach. Nevertheless, as we will see, it has many desirable properties. It also has the considerable advantage of having a very intuitive and *interpretable* definition. Thus, we consider it a first step toward our general theory. This section is less formal and focuses on insights. The rigorous formulation and proofs are provided in Section IV.

Bounded expectation can be motivated by the de minimis risk principle, which is a generally accepted principle [61], [62]. De minimis is also referred to as the principle of ignoring rationally negligible probabilities (RNP). The RNP principle states that we should ignore very small probabilities, say below ϵ . Indeed, this is what we always do, as in anything we do in real life, there is always a chance of a catastrophe, and we ignore this if the probability is small enough.

Although the RNP principle is to some extent accepted, there is no principled way to apply it in decision-making [63]. For example, suppose an agent is offered a gamble in which a coin is being tossed 100 times, and depending on the sequences of heads and tails some rewards are offered to the agent.³ In such a gamble, the probability of any outcome is below the ϵ threshold (the probability of each outcome is $\frac{1}{2^{100}}$); therefore, it is not clear which outcomes the agent should throw away when applying the RNP principle?

This leads us to BE, which is a very simple version of our nonasymptotic metrics for evaluating the decision choices proposed in Section IV. Consider a scenario where an agent is considering one of m possible actions or choices. The random variables that represent the rewards (utilities) of potential actions are represented as X_i , for $i = 1, 2, \dots, m$. The standard approach of maximizing expected utility advises choosing the action with the highest expected value (utility), $E[X_i]$. We now provide the BE metric, denoted by E_ϵ , as a way to measure the value of each action.

The basic idea is very simple: We first identify “extreme values” (outliers) of X_i from the right and left in such a way that the probabilities of such extreme values are in total less than or equal to $\frac{\epsilon}{m}$. Here, ϵ is what we consider the rationally negligible probability, and m is the number of alternative options we are considering, i.e., the number of random variables. The BE of X_i , shown as $E_\epsilon[X_i]$, is then the conditional expected value of X_i given that X_i is not in the outlier region.

If the X_i s are continuous random variables, then the definition can be simplified as follows. First, for each X_i , we identify the values of $x_{i,\min}(\epsilon)$ and $x_{i,\max}(\epsilon)$ in a way that the tail probabilities $P(X_i < x_{i,\min}(\epsilon))$ and $P(X_i > x_{i,\max}(\epsilon))$ are each

³Consider a long table with 2^{100} rows, where the reward for each outcome has been decided possibly by random drawing from a distribution.

equal to $\frac{\epsilon}{2m}$. Figure 1 shows these values for a continuous random variable X with a probability distribution function (PDF) $f_X(x)$. The formal definition of $x_{min}(\epsilon)$ and $x_{max}(\epsilon)$ for a random variable X is as follows:

$$x_{min}(\epsilon) = \inf \left\{ x \in \mathbb{R} : P(X \leq x) \geq \frac{\epsilon}{2m} \right\},$$

$$x_{max}(\epsilon) = \sup \left\{ x \in \mathbb{R} : P(X \geq x) \geq \frac{\epsilon}{2m} \right\}.$$




FIGURE 1. The values of $x_{min}(\epsilon)$ and $x_{max}(\epsilon)$ for a random variable with PDF $f_X(x)$.

Hence, we can state the definition of BE for continuous random variables in the following way. Note that in the definition of bounded expectation below, $x_{i,min}(\epsilon)$ and $x_{i,max}(\epsilon)$ are those values associated with the random variable X_i .

Definition 1 (BE for Continuous Random Variables): To decide between jointly continuous random variables X_i , for $i = 1, 2, \dots, m$, let the A_i s be the events $\{x_{i,min}(\epsilon) \leq X_i \leq x_{i,max}(\epsilon)\}$ and $A = \bigcap_{i=1}^m A_i$. The value of action i associated with X_i based on the BE metric is given by

$$v[X_i] = E_\epsilon[X_i] = E[X_i|A]$$

$$= E \left[X_i | x_{i,min}(\epsilon) \leq X_i \leq x_{i,max}(\epsilon), i = 1, 2, \dots, m \right].$$

The above definition will be extended in a specific way that can be applied to all random variables to ensure that some regularity conditions are satisfied (which will be discussed in Section IV). Nevertheless, the basic idea shown in Figure 1 and the above definition for continuous random variables are sufficient for our discussions in this section.

Note that the BE metric cannot simply be expressed as $E[X_i|a \leq X_i \leq b]$, where a and b are constants. Indeed, the difference is that in BE, the values of a and b depend on the distribution of X_i and are different for each of the X_i s. Additionally, if the X_i s are not independent, the event A_i impacts the E_ϵ for other X_j s.

The intuition behind BE is that we precisely throw out the “outliers” in the distributions of X_i s to focus on the part of the probability space that will happen with a very high probability. Note that, as will be discussed, this does not

mean we are ignoring tail risks. Indeed, we will see that BE can be used to describe and analyze such risks. Moreover, the extensions and generalizations that we propose later do not throw out any part of the probability space.

In the RNP literature, there are discussions on how to choose the value below which we ignore probabilities [64], [65]. This is equivalent to the value of ϵ in BE. The suggested values usually range between 10^{-3} and 10^{-6} .

Next, we will discuss using BE for decision-making. We focus on the insights and discussions and leave the proof for Section IV. Note that if X_i ’s are independent continuous random variables or there is only one $X_i = X$ (the agent is deciding how much to pay for X), we can simply write

$$E_\epsilon[X] = E[X|A] = E \left[X | x_{min}(\epsilon) \leq X \leq x_{max}(\epsilon) \right].$$

A. BE FOR ST. PETERSBURG GAME

Figure 2 shows the BE value of the St. Petersburg game (per play) as a function of n , the number of times the agent is allowed to play the game, for $\epsilon = 0.001$. Specifically, if the agent plays the game n times independently, and $X^{(j)}$ ⁴ shows the reward on the j th play, we can define the sample mean as

$$\bar{X}_n = \frac{1}{n} \sum_{j=1}^n X^{(j)}.$$




FIGURE 2. BE vs. the logarithm of the number of St. Petersburg game iterations, assuming $\epsilon = 10^{-3}$.

The per game value according to BE is given by $E_\epsilon [\bar{X}_n]$. As we see, this value increases as n becomes larger. Indeed, as $n \rightarrow \infty$, $E_\epsilon [\bar{X}_n] \rightarrow \infty$. This is exactly what we expect for a fair valuation. As one plays more, the probability of

⁴In this article, we normally use subscripts to identify random variables that are associated with different actions and we are interested in comparing them, i.e., X_1, X_2, \dots, X_m . Superscripts with parentheses ($X^{(j)}$) are usually used when we are referring to independent and identically distributed (i.i.d.) random variables. The subscripts in parentheses ($X_{(i)}$) are used when we refer to the order statistics of i.i.d. samples from a distribution. Finally, superscripts with brackets ($X^{[n]}$) are used when we refer to sequences of random variables that converge to another random variable.

extreme observations increases, which increases the value of the game.

B. NONLINEARITY OF BE

Note that for BE, we do not necessarily have

$$E_\epsilon[X + Y] = E_\epsilon[X] + E_\epsilon[Y].$$

This is why for the St. Petersburg game, independent repetitions of the game increase the per-game value. Indeed, for independent *right-tailed* random variables X and Y , we often observe

$$E_\epsilon[X + Y] > E_\epsilon[X] + E_\epsilon[Y]$$

This nonlinearity is pronounced for heavy-tailed distributions.

Figure 3 shows the BE for Pareto random variables with infinite expected value. As represented, the BE of the sample mean $E_\epsilon[\bar{X}_n]$ is more than the average of the BEs, i.e.,

$$\begin{aligned} E_\epsilon[\bar{X}_n] &= E_\epsilon\left[\frac{1}{n} \sum_{j=1}^n X^{(j)}\right] \\ &> \frac{1}{n} \sum_{j=1}^n E_\epsilon[X^{(j)}] \\ &= E_\epsilon[X^{(j)}]. \end{aligned}$$




FIGURE 3. BE of the sample mean of the right-tailed Pareto random variables, $E_\epsilon[\bar{X}_n]$, and the BE of $X^{(j)}$'s vs. the number of repetitions n for the case of infinite expected value, assuming $\epsilon = 0.001$.

Furthermore, the BE converges to the actual expectation as the number of experiments increases. Figure 4 shows this fact for the Pareto random variable with finite expected value. In both Figures 3 and 4, ϵ is assumed to be 0.001.




FIGURE 4. BE of the sample mean and the BE of $X^{(j)}$'s of the right-tailed Pareto random variable with vs. the number of repetitions n for the case of finite expected value, assuming $\epsilon = 0.001$.

C. ON THE PROFITABILITY OF THE VENTURE CAPITAL INDUSTRY

Let us apply BE to a concrete example, venture capital investing. A venture capitalist invests in a large number (L) of right-tailed (usually heavy-tailed) options, $X^{(j)}$, $j = 1, 2, \dots, L$. As discussed above, for such right-tailed distributions, assuming $X^{(j)}$'s are independent, we often have

$$E_\epsilon\left[\sum_{j=1}^L X^{(j)}\right] \gg \sum_{j=1}^L E_\epsilon[X^{(j)}].$$

Thus, the aggregate value of the investment portfolio is much larger than the sum of individual items. Each startup company has a very small chance of success. It has a still smaller chance of great success, so its individual valuation is small. However, the aggregate is much more valuable, as predicted by the BE measure. We will explore this topic more deeply in Section V-C after the full theory is developed.

D. ON TAIL RISKS, ROBUSTNESS, AND FRAGILITY

The situation is reversed for left-tailed distributions. For independent random variables X and Y that are right-tailed, we often have

$$E_\epsilon[X + Y] < E_\epsilon[X] + E_\epsilon[Y].$$

Figure 5 shows this relation for left-tailed Pareto random variables with infinite expected value. As shown, the BE for the sum of the left-tailed random variables is less than the sum of the BEs of each random variable. Again, $\epsilon = 0.001$. Similarly, Figure 6 shows this phenomenon for the Pareto random variable with finite expected value.

What does all this mean? The above can describe the situation of accumulation of risks. An agent might be taking risks that are individually acceptable but not acceptable on average. This is related to the issue of robustness and




FIGURE 5. BE of the sample mean of left-tailed Pareto random variables $E_\epsilon[\bar{X}_n]$ and the BE of $X^{(j)}$'s vs. the number of repetitions n for the case of infinite expected value, $E[X^{(j)}] = -\infty$, assuming $\epsilon = 0.001$.




FIGURE 6. BE of the sample mean and the BE of $X^{(j)}$'s of the left-tailed Pareto random variable vs. the number of repetitions n for the case of finite expected value, assuming $\epsilon = 0.001$.

fragility [13], [14], [66]–[69]. This is the reverse of the situation in the venture capital example. Each individual risk is very limited. The probability of a loss for each action might be small. The probability of a large loss is even much smaller. However, the aggregate risk is by far larger than the sum of the risks.

E. LIMITATIONS OF BE: TOWARD A MORE GENERAL FRAMEWORK

The above discussions were intended to show that the BE measure can be considered a simple measure that has several desirable properties: In addition to providing a satisfactory answer to problems such as the St. Petersburg problem, it could provide insights on some practical situations. Nevertheless, it is not perfect (like any other measure). The problem is that it only solves one issue regarding one-time

decisions: the RNP issue. It does not address the rest of the probability space. BE seems to provide a satisfactory answer to the St. Petersburg problem, but let us consider our second example regarding the legal case. For that case, we obtain

$$E_\epsilon[X] = 10, \\ E_\epsilon[Y] = \frac{10.75 - 17.5\epsilon}{1 - \epsilon}.$$

It is easy to verify that if ϵ is small, the result of BE is not very different from what is predicted by the expected utility. Thus, we need a more comprehensive approach that looks at the entire probability space, not just a negligible part.

To better motivate our general framework, let us now look at different views of BE. The BE operation can be thought of as normal expectation in a “distorted” or “modified” probability space, one for which the probabilities of events in A^c are reduced to zero, but the probabilities of events in A are multiplied by $\frac{1}{P(A)}$. The intuition is that we are magnifying the most likely outcomes and shrinking highly improbable outcomes since we are focusing on one-time decisions that cannot be repeated; thus, we do not have the luxury provided by the LLN.

Now, this idea of modifying the probability space to better accommodate the lack of repetition and the non-LLN regime can be developed into a much more general methodology. For example, there is no reason to partition the space into only two subsets. We can partition into more sets and adjust the probability of each part in a specific way. Additionally, BE applies an abrupt change, i.e., reduces some probabilities to zero. We can instead change the probabilities in a smoother way and still enjoy the aforementioned attractive properties of BE as well as many more. All of these are covered under the proposed framework of the change-of-probability measure in the next sections. For example, we see that the general approach provides a more satisfactory answer to the legal case question.

It is clear that this change-of-probability measure operation (probability modification) cannot be arbitrary and must be done in a principled way to ensure that it is consistent with rational decision-making. Therefore, in the next section, we develop a rigorous theory for such change-of-probability measure operations for decision-making and prove their properties. This will lead to the systematic decision-making in non-LLN regimes discussed in Section V.

IV. A GENERAL FRAMEWORK: CHANGE-OF-PROBABILITY MEASURE

In this section, we develop a general framework based on the change-of-the probability measure for probabilistic decision-making. The idea is to list some important properties that such change-of-measure operations must satisfy. Change-of-measure policies that satisfy these properties are called ϵ -consistent. The parameter ϵ is a measure of “distortion” imposed on the probability measures and plays a key role in our analysis (Property 6 below). We then prove some properties (such as convergence to expected utility as the

number of repetitions grows) that all ϵ -consistent policies satisfy. We first focus on the case of *uniform* change-of-measures and then discuss *nonuniform* change-of-measures. BE is proven to be uniform ϵ -consistent.

Later, in Section V, we will focus on two important tasks. First, we will provide a specific method for constructing ϵ -consistent change-of-measure policies using what we call *consistent functions*. Second, and more importantly, we will propose a systematic approach to applying ϵ -consistent policies in probabilistic decision-making.

As our goal is to build a rigorous theory, many of the forthcoming sections are somewhat technical. Readers less interested in the mathematical details can refer to Sections V-A, V-B, and V-C to see a summary of the approach and examples of how it can be used in practice.

A. UNIFORM ϵ -CONSISTENT CHANGE-OF-MEASURE POLICIES

Consider a complete probability space (Ω, \mathcal{F}, P) . The random variables that represent the rewards (utilities) of potential actions are defined on this probability space. Concretely, let $X_i : \Omega \mapsto \mathbb{R}$, for $i = 1, 2, \dots, m$, show the reward associated with the m potential actions that we are considering. It is in general convenient (and not restrictive) if we assume \mathcal{F} is the sigma field generated by all the involved random variables. So we are making this assumption unless stated otherwise. For example, if we are considering a fixed set of random variables X_i , for $i = 1, 2, \dots, m$, we may assume $\mathcal{F} = \sigma(X_1, X_2, \dots, X_m)$, where, $\sigma(X_1, X_2, \dots, X_m)$ is the sigma field generated by X_1, X_2, \dots, X_m .

The goal here is to define a new probability measure Q on (Ω, \mathcal{F}) to be used in evaluating the true value of these actions. More specifically, for a generic random variable X , the value $v[X]$ is obtained by

$$v[X] = \int_{\Omega} X(\omega) dQ(\omega). \quad (1)$$

Intuitively, the new probability measure can be defined in a way to potentially amplify the most likely outcomes while weakening the highly unlikely outcomes. This could be consistent with the nonasymptotic nature of the problem, for example, a one-time high-impact decision. Therefore, our goal is to describe mappings $P \mapsto Q$ that have desirable properties consistent with probabilistic decision-making.

First, we notice that the measure Q must be absolutely continuous with respect to the P , i.e., $Q \ll P$. This is because any event that has zero probability under P must have zero probability under Q . Thus, we can use the Radon-Nikodym theorem [70] to conclude that there exists a unique integrable nonnegative random variable Z , with $E[Z] = 1$, such that

$$Q(B) = E[1_B Z], \quad \text{for any } B \in \mathcal{F},$$

where

$$1_B = \begin{cases} 1 & \omega \in B \\ 0 & \omega \notin B. \end{cases}$$

Here Z is the Radon-Nikodym derivative $Z = \frac{dQ}{dP}$. Therefore, our goal can be equivalently stated as obtaining a mapping

$$\{X_1, X_2, \dots, X_m\} \mapsto Z,$$

that maps any set of random variables on (Ω, \mathcal{F}, P) to an integrable nonnegative random variable Z . Identifying this Z uniquely identifies the measure Q as well as $v[X_i]$ for $i = 1, 2, \dots, m$. Hence, to summarize this change-of-measure operation, we write $\{X_1, X_2, \dots, X_m, P\} \xrightarrow{\text{ch}} \{Z, Q, v[\cdot]\}$, where “ch” stands for the change-of-probability measure.

For example, if we consider continuous random variables X_i , and for $i = 1, 2, \dots, m$, define

$$R_{X_i, \epsilon} = [x_{i, \min}(\epsilon), x_{i, \max}(\epsilon)],$$

$$A_i = \{\omega \in \Omega : X(\omega) \in R_{X_i, \epsilon}\},$$

and

$$A = \bigcap_{i=1}^m A_i,$$

then, the Z associated with the BE measure that was introduced in the previous section is given by

$$Z(\omega) = \begin{cases} \frac{1}{P(A)} & \omega \in A \\ 0 & \text{otherwise.} \end{cases}$$

In general, we can rewrite Equation (1) as

$$v[X] = \int_{\Omega} X(\omega) dQ(\omega)$$

$$= \int_{\Omega} X(\omega) Z(\omega) dP(\omega) = E[X(\omega)Z(\omega)]. \quad (2)$$

Equation (2) provides the two interpretations for our problem: (1) the change-of-measure interpretation ($P \mapsto Q$) and (2) the transformation interpretation, given by

$$Y(\omega) = X(\omega)Z(\omega),$$

and letting

$$v[X] = E[Y(\omega)],$$

where the expectation is computed with respect to the original probability measure. Both interpretations are helpful and help us gain insight.

Obviously, the change-of-measure operation $P \mapsto Q$ (or equivalently, defining Z) cannot be arbitrary and must be done in a principled way to satisfy some required properties. Hence, we now proceed to identify the properties that should be satisfied by Z and the associated Q and $v[X]$. The properties are listed below.

Property 1 (Finiteness⁵):

- 1) If $P(\{\omega \in \Omega : X(\omega) < \infty\}) = 1$, then $v[X] < \infty$.
- 2) If $P(\{\omega \in \Omega : X(\omega) > -\infty\}) = 1$, then $v[X] > -\infty$.

This property simply states that if under all possible scenarios, we receive a finite reward, then the value $v[X]$ must be

⁵The finiteness property is required only for $\epsilon > 0$, where epsilon is discussed in Property 6. More specifically, we require Property 1, only when $\sup\{|P(B) - Q(B)| : B \in \sigma(X_1, X_2, \dots, X_m)\} > 0$.

finite. Although this might seem a very rational assumption, it is worth noting that this property is not satisfied by standard utility theory (e.g., in the case of St. Petersburg problem). Thus, we depart from the standard expected utility approach from the start.

The next property states that if one option always results in a better outcome, then its value must be higher. This is known as the dominance principle in decision theory.

Property 2 (Dominance): If $P(\{\omega \in \Omega : X_1(\omega) \leq X_2(\omega)\}) = 1$, then $v[X_1] \leq v[X_2]$.

Next, we point our attention to a “fairness” property. Specifically, we require that if we re-order X_i s, none of the $v[X_i]$ s should change. Also, the cumulative distribution function (CDF) of X_i s should be sufficient for determining $v[X_i]$ s. Let us write $[X_1, X_2, \dots, X_m] \mapsto [v_1, v_2, \dots, v_m]$ to mean $v[X_i] = v_i$, for $i = 1, 2, \dots, m$.

Property 3 (CDF Sufficiency/Symmetry): consider m random variables $X_i : \Omega \mapsto \mathbb{R}$ for $i = 1, 2, \dots, m$ with the joint cumulative distribution function (CDF) $F_{X_1, X_2, \dots, X_m}(x_1, x_2, \dots, x_m)$, and let $[X_1, X_2, \dots, X_m] \mapsto [v_1, v_2, \dots, v_m]$.

- 1) *The values $v_i = v[X_i]$, for $i = 1, 2, \dots, m$, are uniquely determined by F_{X_1, X_2, \dots, X_m} .*
- 2) *For any permutation $\pi : \{1, 2, \dots, m\} \mapsto \{1, 2, \dots, m\}$, we must have*

$$[X_{\pi(1)}, X_{\pi(2)}, \dots, X_{\pi(m)}] \mapsto [v_{\pi(1)}, v_{\pi(2)}, \dots, v_{\pi(m)}].$$

As a special case of Property 3, let's consider the case where the joint CDF is symmetric. Specifically, we say that the joint CDF is symmetric if for any $(x_1, x_2, \dots, x_m) \in \mathbb{R}^m$ and a permutation $\pi : \{1, 2, \dots, m\} \mapsto \{1, 2, \dots, m\}$, we have

$$\begin{aligned} F_{X_1, X_2, \dots, X_m}(x_1, x_2, \dots, x_m) \\ = F_{X_1, X_2, \dots, X_m}(x_{\pi(1)}, x_{\pi(2)}, \dots, x_{\pi(m)}). \end{aligned}$$

Now, property 3 ensures that if the joint CDF of the X_i s, $i = 1, 2, \dots, m$, is symmetric, we must have

$$v[X_1] = v[X_2] = \dots = v[X_m],$$

which again indicates fairness.

The next property states that if two options can be made arbitrarily close to each other, their values must also be close to each other. In other words, we are ensuring the continuity of the $v[\cdot]$ metric. We say that a sequence of random variables $X^{[n]}$, $n = 1, 2, 3, \dots$ are dominated in absolute value by the random variable Y , if $|X^{[n]}(\omega)| \leq Y(\omega)$, for all $\omega \in \Omega$ and for all $n = 1, 2, 3, \dots$.

Property 4 (Convergence): Let $X_1^{[n]}$, $n = 1, 2, 3, \dots$, be a sequence of random variables on (Ω, \mathcal{F}, P) and $Q^{[n]}$ be the corresponding measure when the $X_1^{[n]}$ s are being compared to some other fixed set of options X_2, \dots, X_m on the same probability space. Specifically, $\{X_1^{[n]}, X_2, \dots, X_m, P\} \xrightarrow{ch} \{Z^{[n]}, Q^{[n]}, v^{[n]}[\cdot]\}$. Suppose

$$X_1^{[n]} \xrightarrow{a.s.} X_1,$$

where “ $\xrightarrow{a.s.}$ ” indicates almost sure convergence (with respect to P), and $\{X_1, X_2, \dots, X_m, P\} \xrightarrow{ch} \{Z, Q, v[\cdot]\}$. Assume that all the $X_1^{[n]}$ s are dominated in absolute value by an integrable (with respect to P and $Q^{[n]}$) random variable Y . We then have

$$\lim_{n \rightarrow \infty} v^{[n]}[X_1^{[n]}] = v[X_1].$$

Note that since $Q^{[n]} \ll P$, almost sure convergence with respect to P also ensures almost sure convergence with respect to all $Q^{[n]}$ s.

Next, note that $v[X]$ is not necessarily a linear operator as we saw regarding the bounded expectation operator. This is in contrast to standard expected utility. Nevertheless, we require a weaker form of linearity for $v[X]$. Specifically, we require:

Property 5 (Weak Linearity): $v[aX + b] = av[X] + b$, for any $a, b \in \mathbb{R}$.

The “ $+b$ ” part simply says that first obtaining X and then obtaining a constant reward b is equivalent to obtaining the reward $X + b$. The “ aX ” part says that if we multiply all the possible outcomes by a factor of a , it makes sense that the whole value is multiplied by a . Note that this has nothing to do with the concept of marginal utility: it does not say that if we are given twice the money, our utility is multiplied by two. It simply says that if under Action 1, we always get twice the utility compared to Action 2, then Action 1 is worth twice Action 2. Indeed, this property is satisfied in standard utility theory.

This property also implies that $v[-X] = -v[X]$. Note that $v[X]$ is how much the option X is worth to an agent. It simply states that if under Action 1, an agent always obtains the negative utility that she obtains under Action 2, then Action 1 must have the negative aggregate value as Action 2. Note that this is not inconsistent with incorporating issues such as loss aversion, as those can be incorporated in the way we define utilities, so if, for example, under Action 1, the agent loses \$100 for sure and under Action 2, she wins \$100 for sure, the utility of Action 1 could be -150 , while the utility of Action 2 could be 100 . For this choice, $v[\text{Action 1}] \neq -v[\text{Action 2}]$. Moreover, to incorporate issues such as risk aversion or risk-seeking, one may use techniques such as risk-weighted expected utility [33] in conjunction with the change-of-measure operation proposed here. Note that Property 5 implies $v[b] = b$ for $b \in \mathbb{R}$.

The next property makes sure we do not distort the probabilities too much. This is crucial, as the agent is basing her decision on the modified probability measure Q , so we would like to make sure that for any event B , the actual probability of that event, i.e., $P(B)$, is within ϵ of its distorted probability. Specifically:

Property 6 (Bounded Distortion): For any event $B \in \sigma(X_1, X_2, \dots, X_m)$, we must have

$$|P(B) - Q(B)| \leq \epsilon.$$

This property essentially says that the total variation distance between the two probability measures P and Q on $\sigma(X_1, X_2, \dots, X_m)$ must be less than or equal to ϵ . The

parameter ϵ plays a key role in the change-of-measure operation. As we will see by increasing the ϵ from zero to positive values and examining the outcomes, we develop a systematic approach to decision-making under non-LLN regimes. It is worth noting that in this context, the interpretation of ϵ is broader than our previous narrow interpretation of BE: here, it refers to the maximum distortion in probabilities. Nevertheless, as we will see, the values of both interpretations coincide in the special case of BE.

The following lemma is useful in constructing bounded distortion change-of-measure operations.

Lemma 1 (Partition Lemma): Let D_1, D_2, \dots, D_m with $P(D_i) > 0$ be sets in \mathcal{F} that form a partition of Ω . Let $\alpha_i \geq 0$ for $i = 1, 2, \dots, m$ be such that

$$\sum_{i=1}^m \alpha_i P(D_i) = 1, \quad \text{and} \quad \sum_{i=1}^m |(1 - \alpha_i)| P(D_i) \leq 2\epsilon.$$

If we define for any $B \in \mathcal{F}$,

$$Q(B) = \sum_{i=1}^m \alpha_i P(B \cap D_i),$$

then Q defines a probability measure on (Ω, \mathcal{F}) , and for any $B \in \mathcal{F}$, we have

$$|P(B) - Q(B)| \leq \epsilon.$$

Proof: First, we note that we have:

$$Q(\Omega) = \sum_{i=1}^m \alpha_i P(D_i) = 1.$$

Let $\delta(P, Q)$ show the total variation distance between P and Q which is given by

$$\delta(P, Q) = \sup_{A \in \mathcal{F}} |P(A) - Q(A)|.$$

Since D_i 's form a partition of Ω , we have

$$\begin{aligned} \delta(P, Q) &= \frac{1}{2} \sum_{i=1}^m |P(D_i) - Q(D_i)| \\ &= \frac{1}{2} \sum_{i=1}^m |(1 - \alpha_i)| P(D_i) \\ &\leq \epsilon. \end{aligned}$$

Thus, for any $B \in \mathcal{F}$, we have

$$|P(B) - Q(B)| \leq \epsilon.$$

Definition 2 (Uniform ϵ -Consistent Policies): Consider a probability space (Ω, \mathcal{F}, P) and a mapping rule that maps any set of random variables on (Ω, \mathcal{F}, P) to an integrable nonnegative random variable Z , with $E[Z] = 1$. Assume Q and $v[\cdot]$ are the associated measure and the value function. We say that this change-of-measure operation is a uniform ϵ -consistent policy if it satisfies properties 1 through 6.

When we simply say ϵ -consistent, we mean uniform ϵ -consistent. We now proceed to prove some properties of ϵ -consistent change-of-measure policies. We say that the random variable X is symmetric around μ if $2\mu - X$ has the

same distribution as X . Our first theorem considers the case where we are evaluating $v[X]$ for a single ($m = 1$) symmetric random variable.

Theorem 1: Let $v[\cdot]$ be associated with an ϵ -consistent change-of-measure policy applied to a single random variable X . If X is symmetric around μ , then $v[X] = \mu$.

Proof: We have

$$\begin{aligned} v[X] &= v[2\mu - X] \\ &= 2\mu - v[X]. \end{aligned}$$

The first equality is true by Property 3 and the assumption that X is symmetric around μ . The second equality is true by Property 5. We conclude that $v[X] = \mu$. \blacksquare

Note: It is crucial to note that the proposed framework here addresses a single issue: the non-LLN nature of some decision problems. Issues such as attitudes toward risk can be further combined with the proposed method to obtain a more comprehensive view. For example, in financial investment, it is common to prefer options with a lower variance when the expected values are the same even when the distributions are symmetric around the mean.

We now turn our attention to the case where an agent is able to independently repeat an action several times; here, \bar{X} indicates the average reward.

$$\bar{X}_n = \frac{X^{(1)} + X^{(2)} + \dots + X^{(n)}}{n}.$$

In this case, we show that $v[\bar{X}_n]$ converges to $E[X]$. This means that as the number of repetitions of an action grows, we approach the expected utility theory.

Theorem 2 (Limit Theorem): Let $X^{(1)}, X^{(2)}, \dots, X^{(n)}$ be i.i.d. with expected value $E[X^{(i)}] = \mu < \infty$. Let $v[\cdot]$ be associated with an ϵ -consistent change-of-measure policy. Assume \bar{X}_n 's are dominated in absolute value by an integrable random variable Y . We have

$$\lim_{n \rightarrow \infty} v[\bar{X}_n] = \mu.$$

Proof: Since the $X^{(i)}$'s are i.i.d. with finite mean, we can use the strong law of large numbers and conclude that

$$\bar{X}_n \xrightarrow{a.s.} \mu.$$

Since \bar{X}_n 's are dominated in absolute value by an integrable random variable Y , by Property 4, we conclude that

$$\begin{aligned} \lim_{n \rightarrow \infty} v[\bar{X}_n] &= v[\mu] \\ &= \mu. \end{aligned}$$

The last equality is ensured by Property 5. \blacksquare

We can often provide a stronger characterization under some regularity conditions. Specifically, suppose that $X^{(i)}$'s have a finite variance: $0 < \text{Var}(X^{(i)}) = \sigma^2 < \infty$. We can apply the central limit theorem (CLT) and conclude that the random variables

$$W_n = \frac{\bar{X}_n - \mu}{\sigma/\sqrt{n}},$$

converge in distribution to the standard normal random variable $W \sim N(0, 1)$ as n goes to infinity. In such cases it is often the case that $v[W_n]$ also converges to $v[W]$ under some regularity conditions. Since by Theorem 1, $v[W] = 0$, we may conclude

$$v \left[\frac{\bar{X}_n - \mu}{\sigma/\sqrt{n}} \right] \rightarrow 0.$$

By Property 5, we may conclude

$$\frac{\sqrt{n}}{\sigma} (v[\bar{X}_n] - \mu) \rightarrow 0.$$

B. BOUNDED EXPECTATION REVISITED

We now formally define BE and show that it provides an ϵ -consistent change-of-measure. Therefore, it enjoys the properties discussed above. There are two equivalent ways to define BE: One is more suitable for simulations; we call it the *operational definition*. The other is more suitable for direct calculations; we call it the *computational definition*. For simplicity, let us first assume that the m random variables associated with different options are independent. The definition then will simply be extended to the case when they are not independent. Let us start with the operational definition, which is more intuitive.

Let X be the random variable indicating the reward (utility) of the action. Assume $X(\omega) < \infty$ for all $\omega \in \Omega$. Generate N i.i.d. random variables from the distribution $F_X(x)$, order them from smallest to largest, and denote the resulting sequence of random variables as

$$X_{(1)}, X_{(2)}, \dots, X_{(N)}.$$

In other words, $X_{(1)}, X_{(2)}, \dots, X_{(N)}$ is the order statistic of the random sample. Define the “normal” set I_N as

$$I_N = \left\{ \lfloor \frac{N\epsilon}{2m} \rfloor, \lfloor \frac{N\epsilon}{2m} \rfloor + 1, \lfloor \frac{N\epsilon}{2m} \rfloor + 2, \dots, N - \lfloor \frac{N\epsilon}{2m} \rfloor \right\}.$$

Accordingly, the outlier set is defined as $\{1, 2, \dots, N\} - I_N$. Figure 7 shows this sample division for a Pareto random variable. In this figure, $\epsilon = 0.1$ for the sake of representation.

Define the random variables $v_N[X]$ as

$$v_N[X] = \frac{1}{|I_N|} \sum_{i \in I_N} X_{(i)}. \quad (3)$$

The $v_N[X]$ s converge almost surely to a finite limit, which we call the bounded expectation of X :

$$E_\epsilon[X] = \lim_{N \rightarrow \infty} v_N[X]. \quad (4)$$

To see that the limit exists and is finite, we actually derive the limit that gives us the computational definition of BE. Specifically, by applying the law of large numbers, we obtain

$$E_\epsilon[X] = \left(P(X \leq x_{min}(\epsilon)) - \frac{\epsilon}{2m} \right) \frac{x_{min}(\epsilon)}{1 - \epsilon} + \left(P(X \geq x_{max}(\epsilon)) - \frac{\epsilon}{2m} \right) \frac{x_{max}(\epsilon)}{1 - \epsilon}$$




FIGURE 7. Representation of the operational definition: Order statistics of the Pareto random variable. $\lfloor \frac{N\epsilon}{2m} \rfloor$ of the first and last samples are put in the outlier set. $\epsilon = 0.1$ (Relatively large ϵ has been chosen to facilitate the representation).

$$+ \left(P(x_{min}(\epsilon) < X < x_{max}(\epsilon)) \right) \times \frac{E[X|x_{min}(\epsilon) < X < x_{max}(\epsilon)]}{1 - \epsilon}. \quad (5)$$

Here,

$$x_{min}(\epsilon) = \inf \left\{ x \in \mathbb{R} : P(X \leq x) \geq \frac{\epsilon}{2m} \right\}, \quad x_{max}(\epsilon) = \sup \left\{ x \in \mathbb{R} : P(X \geq x) \geq \frac{\epsilon}{2m} \right\}. \quad (6)$$

To consider the case where the X_i s are not necessarily independent, we can proceed similarly. Specifically, we consider the vector X :

$$X = \begin{bmatrix} X_1 \\ X_2 \\ \vdots \\ X_m \end{bmatrix}.$$

with joint CDF

$$F_X(x) = F_{X_1, X_2, \dots, X_m}(x_1, x_2, \dots, x_m).$$

Generate N i.i.d. random variables from the distribution F_X . For each $i \in \{1, 2, \dots, m\}$, order the obtained vectors in terms of the value of X_i (for equal values, the ordering is random). Then, label the vectors as outliers as before. Finally, all the vectors that have been labeled as outliers at least once, are removed. Note that by the union bound, the outlier set has at most $N\epsilon$ elements. The remaining collection of samples will give us the normal set, and $E_\epsilon[X]$ can be computed as before using Equations (3) and (4).

We now state and prove the main theorem regarding BE.

Theorem 3: BE is an ϵ -consistent change-of-measure operation.

Proof:

Property 1: If $P(\omega \in \Omega : X(\omega) < \infty) = 1$, then for any $\epsilon > 0$, we must have $x_{max}(\epsilon) < \infty$. Similarly,

if $P(\{\omega \in \Omega : X(\omega) > -\infty\}) = 1$, then we must have $x_{min}(\epsilon) > -\infty$, for any $\epsilon > 0$. Now note that for any random variable X_i , $i = 1, 2, \dots, m$, we have

$$x_{i,min}(\epsilon) \leq v[X_i] \leq x_{i,max}(\epsilon),$$

which finishes the proof.

Property 2: Let $\{X_1, X_2, P\} \xrightarrow{ch} \{Z, Q, v[\cdot]\}$. Additionally, let $\Gamma = \{\omega \in \Omega : X_1(\omega) > X_2(\omega)\}$. Since $P(\Gamma) = 0$ and $Q \ll P$, we conclude that $Q(\Gamma) = 0$. Then,

$$\begin{aligned} v[X_1] &= \int_{\Omega} X_1(\omega) dQ(\omega) \\ &\stackrel{(a)}{=} \int_{\Omega - \Gamma} X_1(\omega) dQ(\omega) \\ &\stackrel{(b)}{\leq} \int_{\Omega - \Gamma} X_2(\omega) dQ(\omega) \\ &= \int_{\Omega} X_2(\omega) dQ(\omega) \\ &= v[X_2], \end{aligned}$$

where (a) comes from the fact that $Q(\Gamma) = 0$ and (b) holds according to the assumption that $X_1(\omega) \leq X_2(\omega)$ for all $\omega \in \Omega - \Gamma$.

Property 3: Since BE is computed and uniquely determined by the CDF, and is symmetric with respect to the ordering of random variables, this property is satisfied.

Property 4: For simplicity, let $m = 1$. We first show that it suffices to prove the statement for continuous random variables. Let U be a random variable uniformly distributed in $[-1, 1]$, and $\delta > 0$. Define

$$Y^{[n]} = X^{[n]} + \delta U,$$

and

$$Y = X + \delta U.$$

Note that random variables Y and $Y^{[n]}$, $n \in \mathbb{N}$, are continuous and $Y^{[n]} \xrightarrow{a.s.} Y$. Also, we have:

$$\begin{aligned} |E_{\epsilon}[Y] - E_{\epsilon}[X]| &\leq \delta, \\ |E_{\epsilon}[Y^{[n]}] - E_{\epsilon}[X^{[n]}]| &\leq \delta. \end{aligned}$$

Note that $\delta > 0$ can be chosen arbitrarily small. Therefore, if we have $E_{\epsilon}[Y^{[n]}] \rightarrow E_{\epsilon}[Y]$, then we can conclude that $E_{\epsilon}[X^{[n]}] \rightarrow E_{\epsilon}[X]$. Hence, it suffices to prove the statement for continuous random variables $X^{[n]}$ and X . Let $a = x_{min}(\epsilon)$, $b = x_{max}(\epsilon)$, $a_n = x_{min}^{[n]}$, and $b_n = x_{max}^{[n]}$. We then have

$$\begin{aligned} E_{\epsilon}[X] &= E[X | x_{min}(\epsilon) < X < x_{max}(\epsilon)] \\ &= \frac{1}{1 - \epsilon} E[X 1_{[a,b]}(X)], \\ E_{\epsilon}[X^{[n]}] &= E[X^{[n]} | x_{min}^{[n]}(\epsilon) < X^{[n]} < x_{max}^{[n]}(\epsilon)] \\ &= \frac{1}{1 - \epsilon} E[X^{[n]} 1_{[a_n,b_n]}(X^{[n]})]. \end{aligned}$$

Define

$$h_{\beta}(x) = \begin{cases} 0 & x < a - \beta \\ 1 + \frac{1}{\beta}(x - a) & a - \beta \leq x \leq a \\ 1 & x > a, \end{cases}$$

where $\beta > 0$. We have $P(X^{[n]} \geq a) \leq E[h_{\beta}(X^{[n]})]$. Furthermore, since $h_{\beta}(x)$ is a continuous function, we have

$$\lim_{n \rightarrow \infty} E[h_{\beta}(X^{[n]})] = E[h_{\beta}(X)],$$

and

$$\lim_{\beta \rightarrow 0} E[h_{\beta}(X)] = P(X \geq a) = 1 - \frac{\epsilon}{2},$$

from which we conclude that

$$\limsup_{n \rightarrow \infty} P(X^{[n]} \geq a) \leq 1 - \frac{\epsilon}{2}.$$

Next, we have

$$\begin{aligned} P(X^{[n]} \geq a - \beta) &\stackrel{n \rightarrow \infty}{\geq} \frac{E[h_{\beta}(X^{[n]})]}{E[h_{\beta}(X)]} \\ &\stackrel{(b)}{>} \frac{P(X \geq a)}{1 - \frac{\epsilon}{2}} \\ &= P(X^{[n]} \geq a_n), \end{aligned}$$

where (b) results from the definition of $x_{min}(\epsilon)$ and the assumption that X is continuous. For large enough n , we conclude that

$$P(X^{[n]} \geq a - \beta) > 1 - \frac{\epsilon}{2},$$

which results in $a_n \geq a - \beta$. Similarly, $b_n \leq b - \beta$. Now, since $X^{[n]} \xrightarrow{a.s.} X$, we have

$$X^{[n]} 1_{[a-\beta,b+\beta]}(X^{[n]}) \rightarrow X 1_{[a-\beta,b+\beta]}(X).$$

Applying DCT, we conclude that

$$\begin{aligned} E[X^{[n]} 1_{[a-\beta,b+\beta]}(X^{[n]})] &\stackrel{n \rightarrow \infty}{\rightarrow} E[X 1_{[a-\beta,b+\beta]}(X)] \\ &\stackrel{\beta \rightarrow 0}{\rightarrow} E[X 1_{[a,b]}(X)]. \end{aligned}$$

Thus,

$$\lim_{\beta \rightarrow 0} \lim_{n \rightarrow \infty} E[X^{[n]} 1_{[a-\beta,b+\beta]}(X^{[n]})] = E[X 1_{[a,b]}(X)]. \quad (7)$$

Therefore, to finish the proof, it suffices to show that

$$\begin{aligned} \lim_{n \rightarrow \infty} E[X^{[n]} 1_{[a_n,b_n]}(X^{[n]})] &= \lim_{\beta \rightarrow 0} \lim_{n \rightarrow \infty} E[X^{[n]} 1_{[a-\beta,b+\beta]}(X^{[n]})]. \end{aligned}$$

We have

$$\begin{aligned} P(X^{[n]} \geq a - \beta) &= P(X^{[n]} + \beta \geq a) \\ &\leq E[h_{\gamma}(X^{[n]} + \beta)] \\ &\stackrel{n \rightarrow \infty}{\rightarrow} E[h_{\gamma}(X + \beta)] \\ &\stackrel{\gamma \rightarrow 0}{\rightarrow} P(X + \beta \geq a) \\ &= P(X \geq a - \beta). \end{aligned}$$

Now, we conclude that

$$\limsup_{n \rightarrow \infty} P(X^{[n]} \geq a - \beta) \leq P(X \geq a - \beta).$$

Therefore,

$$\begin{aligned} \liminf_{n \rightarrow \infty} P(X^{[n]} < a - \beta) &\geq P(X < a - \beta) \\ &= P(X < a) - P(a - \beta < X < a) \\ &= \frac{\epsilon}{2} - c(\beta), \end{aligned}$$

where $c(\beta) \rightarrow 0$ as $\beta \rightarrow 0$. Thus, we have

$$\liminf_{n \rightarrow \infty} P(X^{[n]} < a - \beta) \geq \frac{\epsilon}{2} - c(\beta),$$

and

$$P(X^{[n]} < a_n) = \frac{\epsilon}{2}.$$

Now, we conclude that

$$\limsup_{n \rightarrow \infty} P(a - \beta < X^{[n]} < a_n) \leq c(\beta),$$

Hence,

$$\begin{aligned} \limsup_{n \rightarrow \infty} |E[X^{[n]} 1_{[a-\beta, a_n]}(X^{[n]})]| \\ \leq c(\beta) \max(|a - \beta|, |a_n|) \rightarrow 0, \quad \text{as } \beta \rightarrow 0. \end{aligned}$$

Similarly,

$$\limsup_{n \rightarrow \infty} |E[X^{[n]} 1_{[b_n, b+\beta]}(X^{[n]})]| \leq d(\beta) \rightarrow 0, \quad \text{as } \beta \rightarrow 0.$$

We conclude that

$$\begin{aligned} \lim_{\beta \rightarrow 0} \lim_{n \rightarrow \infty} E[X^{[n]} 1_{[a-\beta, b+\beta]}(X^{[n]})] \\ = \lim_{n \rightarrow \infty} E[X^{[n]} 1_{[a_n, b_n]}(X^{[n]})]. \end{aligned}$$

Finally, using (7), we have

$$\lim_{n \rightarrow \infty} E_\epsilon[X^{[n]}] = E_\epsilon[X].$$

Property 5: Consider the operational definition of BE, i.e., Equation 3. Note that the ordering of $Y_{(i)}$ s are exactly the same as that of the $X_{(i)}$ s for $a > 0$ and completely reversed for $a < 0$. The additional $+b$ is added to all $X_{(i)}$ s, so it will appear as a $+b$ in the computation of $v[Y]$.

Property 6: It can be concluded from the application of Lemma 1. ■

C. NONUNIFORM ϵ -CONSISTENT CHANGE-OF-MEASURE POLICIES

It is worth noting that to prove Theorems 1 and 2, we did not specifically use the fact that all random variables go through the same change-of-measure operation. In other words, as long as Properties 1 through 6 are satisfied, we can use Theorems 1 and 2. Therefore, we can actually construct change-of-measure policies as

$$\{X_1, X_2, \dots, X_m, P\} \xrightarrow{ch} \{Z_{X_i}, Q_{X_i}\},$$

such that for any $X \in \{X_1, X_2, \dots, X_m\}$, its value $v[X]$ is given by

$$v[X] = v_X[X] = \int_{\Omega} X(\omega) dQ_X(\omega).$$

We refer to such policies as *nonuniform* change-of-measure policies. Why might we want to consider nonuniform policies? The answer is that they give us more flexibility in defining the appropriate change-of-measures. Specifically, one way to define nonuniform change-of-measure policies, is to act as if $m = 1$ and apply the change-of-measure operation for each X_i separately to obtain $v[X_i]$. The big advantage here is that we do not need to deal with the way the random variables might be dependent on each other. One might argue that, at the end of the day, we are choosing one of the options and all we care is the distribution of the resulting utility. In other words, we might be less concerned with the way the potential options are correlated as we will only be choosing one of them.

Of course, we need to be specifically careful and ensure that the fairness properties, such as dominance property and CDF symmetry property, are satisfied; otherwise, we might be making an unfair comparison. Properties 1 through 5 do not change; however, we provide a slightly modified version of Property 6 to make it suitable for nonuniform change-of-measure policies:

Property 6-b (Bounded Distortion for Nonuniform Policies): For any $i \in \{1, 2, \dots, m\}$ and any event $B \in \sigma(X_i)$, we must have

$$|P(B) - Q_{X_i}(B)| \leq \epsilon.$$

Definition 3 (Nonuniform ϵ -Consistent Policies): Consider a probability space (Ω, \mathcal{F}, P) and a mapping rule that maps any set of random variables on (Ω, \mathcal{F}, P) to a set of integrable nonnegative random variables Z_{X_i} with $E[Z_{X_i}] = 1$. Specifically, we write

$$\{X_1, X_2, \dots, X_m, P\} \xrightarrow{ch} \{Z_{X_i}, Q_{X_i}\},$$

such that for any $X \in \{X_1, X_2, \dots, X_m\}$, its value $v[X]$ is given by

$$\begin{aligned} v[X] = v_X[X] &= \int_{\Omega} X(\omega) dQ_X(\omega) \\ &= \int_{\Omega} X(\omega) Z_X(\omega) dP(\omega). \end{aligned}$$

We say that this change-of-measure operation is a nonuniform ϵ -consistent policy if it satisfies Properties 1 through 5 as well as Property 6-b.

Note that as we discussed above, nonuniform policies could be specially easy to work with when the change of measure operation is performed separately for each random variable. In such cases, it suffices to provide the mapping

$$(\Omega, \sigma(X_i), P) \mapsto (\Omega, \sigma(X_i), Q_{X_i}),$$

for each X_i . That is, for the i th random variable, we restrict our attention to the space $(\Omega, \sigma(X_i))$.

V. SYSTEMATIC APPROACH TO CONSTRUCTING ϵ -CONSISTENT POLICIES

In this section, we focus on two important tasks. First, we will provide a specific method for constructing ϵ -consistent change-of-measure policies using *consistent functions*. Second, and more importantly, we will propose a systematic approach to applying ϵ -consistent policies in probabilistic decision-making. The approach is based on sweeping the parameter ϵ from zero to large values and looking at how preferences change in this process. The idea is to obtain a holistic view of the problem taking into account the complexities involved in decision-making under non-LLN regimes.

As discussed before, BE has some limitations. The BE operation divides the probability space into two parts: the “normal” part and the “outlier” part, which has a probability smaller than ϵ . The BE then completely eliminates the outlier part of the probability space. This is not necessary; the only thing necessary is to weaken that part enough, so that Property 1 is satisfied. Additionally, there is no need to stop at two divisions. We can simply divide the space into more parts. Using the partition lemma (Lemma 1), we can simply construct an ϵ -consistent change-of-measure operation.

In fact, instead of dividing the space, we can apply a *smooth* change-of-measure operation. This seems to have some advantages over the partitioning approach.

Remember that for any random variable X , the tail function $\bar{F}_X(\cdot)$ is given by

$$\bar{F}_X(x) = P(X > x), \quad x \in \mathbb{R}.$$

As before, suppose that we are interested in comparing $X_i : \Omega \mapsto \mathbb{R}$, for $i = 1, 2, \dots, m$. For simplicity, let us adopt the following notation for the CDF and tail function of X_i :

$$F_i(\cdot) = F_{X_i}(\cdot), \quad \bar{F}_i(\cdot) = \bar{F}_{X_i}(\cdot).$$

The key to our method lies in what we call *consistent* functions.

Definition 4 (Consistent Functions): We say that a function $g : [0, 1] \mapsto [0, 1]$ is a consistent function with respect to random variables X_1, X_2, \dots, X_m if all of the following conditions are satisfied:

- 1) g is continuous and increasing, and $g(0) = 0, g(1) = 1$.
- 2) (Lipschitz continuity) There exists $c_g \in \mathbb{R}$ such that $|g(x) - g(y)| \leq c_g |x - y|$ for all $x \in [0, 1]$ and $y \in [0, 1]$.
- 3) On the interval $\left[0, \frac{1}{2}\right]$, $g(\cdot)$ is convex and we have $g(x) \leq x$.
- 4) On the interval $\left[\frac{1}{2}, 1\right]$, $g(\cdot)$ is concave and we have $g(x) \geq x$.
- 5) For all $x \in [0, 1]$,

$$g(x) + g(1 - x) = 1.$$

- 6) For each $i = 1, 2, \dots, m$, there are constants c_i and c'_i in \mathbb{R} such that

$$\int_{c_i}^{\infty} g(\bar{F}_i(x)) dx < \infty, \quad \int_{-\infty}^{c'_i} g(F_i(x)) dx < \infty.$$

It is easy to construct consistent functions, and indeed, there are infinitely many of them for any set of random variables, as we will see. Our main theorem here is the following.

Theorem 4: Let $g : [0, 1] \mapsto [0, 1]$ be a consistent function with respect to random variables X_1, X_2, \dots, X_m . For any $i \in \{1, 2, \dots, m\}$ and $x \in \mathbb{R}$, define

$$Q_{X_i}(X_i > x) = g(P(X_i > x)).$$

Then, $P \mapsto Q_X$ is a nonuniform ϵ -consistent change-of-measure policy, where

$$\epsilon = 2 \sup_{x \in [0, 1]} |g(x) - x|. \quad (8)$$

Intuitively, the consistent-function properties and the change-of-measure operations in Theorem 4 are chosen to weaken the outliers and strengthen the typical outcomes. This is again consistent with our intuition that in a one-shot decision, it makes sense to focus more on the more likely outcomes.

Note: The actual value of ϵ for a specific set of random variables could be smaller than what is stated in Theorem 4. The value in the theorem is chosen so that a general statement can be made. It is easy to verify that, for $m = 1$ and non-atomic probability spaces, BE is a special case of this general policy with

$$g(x) = \begin{cases} 0 & x < \epsilon/2 \\ \frac{x - \epsilon/2}{1 - \frac{\epsilon}{2}} & \epsilon/2 \leq x \leq 1 - \frac{\epsilon}{2} \\ 1 & x > 1 - \frac{\epsilon}{2}. \end{cases}$$

Note that Condition 6 in Definition 4 is very easy to satisfy. All we need is to make sure that the function $g(x)$ becomes relatively flat at $x = 0$ and $x = 1$. One easy way to satisfy it is to apply the BE truncation using a very small value of $\epsilon' = 10^{-4}$. Note that this should be much smaller than the overall ϵ for g . The truncation can be done in a way that continuity is satisfied. Nevertheless, as ϵ' is very small, continuity at that point has a negligible practical impact.

Proof of Theorem 4: For simplicity, for $i \in \{1, 2, \dots, m\}$ we write

$$Q_i = Q_{X_i}.$$

First, note that since g is Lipschitz continuous, we have for any $A \in \sigma(X_i)$,

$$Q_i(A) \leq c_g P(A),$$

so we have $Q_i \ll P$, for $i = 1, 2, \dots, m$.

Property 1: Property 1 is guaranteed due to Condition 5. Specifically,

- 1) $\int_{c_i}^{\infty} g(\bar{F}_i(x)) dx < \infty$ guarantees that $\int_{c_i}^{\infty} Q_i(X_i > x) dx < \infty$.
- 2) $\int_{-\infty}^{c'_i} g(F_i(x)) dx < \infty$ guarantees that $\int_{-\infty}^{c'_i} Q_i(X_i \leq x) dx < \infty$.

Using (1) and (2), we conclude

$$v[X_i] = \int_{\Omega} X_i(\omega) dQ_i(\omega) < \infty.$$

Property 2: If $P(\{\omega \in \Omega : X_1(\omega) \leq X_2(\omega)\}) = 1$, then $P(X_1 > x) \leq P(X_2 > x)$; therefore, for all $x \in \mathbb{R}$, we have

$$\begin{aligned} Q_1(X_1 > x) &= g(P(X_1 > x)) \\ &\stackrel{(a)}{\leq} g(P(X_2 > x)) \\ &= Q_2(X_2 > x), \end{aligned}$$

where (a) results since g is increasing. Therefore, we have $Q_1(X_1 > x) \leq Q_2(X_2 > x)$ for all $x \in \mathbb{R}$. Thus, $v[X_1] \leq v[X_2]$.

Property 3: Since $v[X_i]$ s are uniquely determined by the joint CDF, and the operation is symmetric with respect to the ordering of random variables, this property is satisfied.

Property 4: For simplicity, assume that the random variables are nonnegative. Let

$$\begin{aligned} h(x) &= Q_{X_1}(X_1 > x) = g(P(X_1 > x)), \\ h_n(x) &= Q_{X_1^{[n]}}(X_1^{[n]} > x) = g(P(X_1^{[n]} > x)). \end{aligned}$$

Let Y be the dominating random variable, and

$$u(x) = g(P(Y > x)).$$

For all $x \geq 0$, we have

$$|h_n(x)| \leq u(x).$$

Now note that

$$\begin{aligned} v[X_1] &= \int_0^{\infty} h(x) dx, \\ v[X_1^{[n]}] &= \int_0^{\infty} h_n(x) dx. \end{aligned}$$

Since $|h_n(x)| \leq u(x)$ and $\int_0^{\infty} u(x) dx < \infty$ (Since Y is integrable), we can apply DCT to conclude

$$\lim_{n \rightarrow \infty} \int_0^{\infty} h_n(x) dx = \int_0^{\infty} h(x) dx = v[X_1],$$

which completes the proof.

Property 5: Let $Y = X + b$. Let Q_1 be the measure associated with X and Q_2 be the measure associated with Y . We have

$$\begin{aligned} Q_2(Y > y) &= g(P(Y > y)) = g(P(X + b > y)) \\ &= g(P(X > y - b)) = Q_1(X > y - b) \\ &= Q_1(X + b > y). \end{aligned}$$

Thus,⁶

$$\begin{aligned} v[Y] &= E_{Q_2}[Y] \\ &= E_{Q_1}[X + b] \\ &= E_{Q_1}[X] + b \\ &= v[X] + b. \end{aligned}$$

⁶Here, $E_Q[\cdot]$ shows the expected value with respect to measure Q .

Now let $Y = aX$ and $a > 0$. Hence,

$$\begin{aligned} Q_2(Y > y) &= g(P(Y > y)) \\ &= g(P(aX > y)) \\ &= g(P(X > \frac{y}{a})) \\ &= Q_1(X > \frac{y}{a}) \\ &= Q_1(aX > y). \end{aligned}$$

Therefore:

$$v[Y] = E_{Q_2}[Y] = E_{Q_1}[aX] = aE_{Q_1}[X] = av[X].$$

Similarly, for $a < 0$, we have:

$$\begin{aligned} Q_2(Y > y) &= g(P(Y > y)) \\ &= g(P(aX > y)) \\ &= g(P(X < \frac{y}{a})) \\ &\stackrel{(a)}{=} Q_1(X < \frac{y}{a}) \\ &= Q_1(aX > y), \end{aligned}$$

where (a) results since $g(x) + g(1 - x) = 1$. Therefore,

$$v[Y] = E_{Q_2}[Y] = E_{Q_1}[aX] = aE_{Q_1}[X] = av[X].$$

Property 6-b: Let $\delta_i(P, Q_i)$ show the total variation distance between P and Q_i (measured on $\sigma(X_i)$), given by

$$\delta_i(P, Q_i) = \sup_{A \in \sigma(X_i)} |P(A) - Q_i(A)|.$$

Due to Conditions of Definition 4 (e.g., convexity/concavity of g), it is easy to see that the supremum is obtained by some $B = \{\omega : \theta_1 \leq X_i(\omega) \leq \theta_2\}$. We can then say that for any $i \in \{1, 2, \dots, m\}$ and any event $B \in \sigma(X_i)$, we have

$$\begin{aligned} |P(B) - Q_i(B)| &\leq \sup_{\{\theta : P(X_i \geq \theta) \leq \frac{1}{2}\}} |P(X_i \geq \theta) - Q_i(X_i \geq \theta)| \\ &\quad + \sup_{\{\theta : P(X_i \geq \theta) \geq \frac{1}{2}\}} |P(X_i \geq \theta) - Q_i(X_i \geq \theta)| \\ &\leq \sup_{x \in [0, \frac{1}{2}]} |x - g(x)| + \sup_{x \in [\frac{1}{2}, 1]} |g(x) - x| \\ &\leq 2 \sup_{x \in [0, 1]} |x - g(x)|. \end{aligned}$$

A. A SYSTEMATIC APPROACH TO DECISION-MAKING IN NON-LLN REGIMES

Here, we propose a systematic approach for decision-making in non-LLN regimes. We start by picking an ϵ -consistent policy such as the method described in the previous section using consistent functions. Note that if we let $\epsilon = 0$, we obtain the same results derived from expected utility theory. In general, let $i(\epsilon)$ be the preferred option for a specific ϵ . As we then increase ϵ , we take note of the possible changes in $i(\epsilon)$. Let ϵ^* be the value of ϵ where the first change occurs in $i(\epsilon)$, i.e., $i(\epsilon^*) \neq i(0)$. The key insights are as follows:

- 1) The larger the value of ϵ^* , the more stable is the choice made by the expected utility ($i(0)$). That is, it is more likely that the expected utility is suggesting a good option.
- 2) On the other hand, if the value of ϵ^* is small, this is a high indication that $i(\epsilon^*)$ might be the best choice.

As the problem of decision-making under non-LLN regimes is multifaceted and most likely a simple narrow approach will not be enough, the proposed method above, where we look at how the preferences change as ϵ changes, seems to be a step in the right direction. An interesting question for further research seems to be finding guidelines on the choice of the threshold value of ϵ^* at which $i(\epsilon^*)$ becomes the preferred option. As a very rough rule of thumb, one might suggest $\epsilon^* < 0.05$ might be used as the threshold.

B. EXAMPLE OF THE APPLICATION OF THE CHANGE-OF-MEASURE OPERATION USING CONSISTENT FUNCTIONS

To clearly present the proposed method, let us revisit the problems that we introduced in the motivation section. We choose the following consistent g function:

$$g(x) = \begin{cases} 0 & x \leq 10^{-4} \\ 2^\alpha x^{1+\alpha} & 10^{-4} \leq x \leq 0.5 \\ 1 - g(1 - x) & x \geq 0.5. \end{cases}$$

Note that technically, the above function is discontinuous at 10^{-4} and $1 - 10^{-4}$. However, since the discontinuity jump is so small, it does not have any practical impact on our calculation. Nevertheless, one may easily make the function fully continuous by replacing the jump with a smooth curve. For $\alpha = 0$, we have $g(x) = x$, so we obtain the standard expected utility, and $\epsilon = 0$. As we increase α , ϵ increases. Therefore, for any α , we obtain a corresponding value for ϵ . Thus, to apply our systematic approach, it suffices to increase α gradually and compute the corresponding value of ϵ as outlined below. Figure 8 shows this $g(x)$ for different α values.

If we have a random variable X and are interested in evaluating its value, i.e., $v[X]$, we can proceed as follows. Specifically, if the random variable X is discrete and bounded from the left, we can simplify the change-of-measure operation in Theorem 4 in the following way. Suppose $\{x_1, x_2, \dots\}$ are potential values of X in an ordered way, i.e.,

$$x_1 < x_2 < x_3 \dots$$

Let $p_i = P(X = x_i)$. Then, we obtain the changed probabilities, q_k , for $k = 1, 2, \dots$, as below

$$q_k = Q(X = x_k) = g\left(\sum_{i=k}^{\infty} p_i\right) - g\left(\sum_{i=k+1}^{\infty} p_i\right).$$

If the range is finite, i.e.,

$$x_1 < x_2 < x_3 \dots < x_r,$$




FIGURE 8. A representation of $g(x)$ with different α values.

then, for $k = 1, 2, \dots, r - 1$, we obtain

$$q_k = Q(X = x_k) = g\left(\sum_{i=k}^r p_i\right) - g\left(\sum_{i=k+1}^r p_i\right),$$

and

$$q_r = Q(X = x_r) = g(p_r).$$

The value of ϵ can be obtained using the total variation distance between P and Q which in this case simplifies to

$$\epsilon = \frac{1}{2} \sum_{i=1}^r |p_i - q_i|.$$

The value of X is then obtained as

$$v[X] = v_\epsilon[X] = \sum_k x_k q_k.$$

Algorithm 1 represents the procedure for calculating $v[X]$ and ϵ .

We now can use the above to revisit the problems that we introduced in the motivation section.

Let us first consider the St. Petersburg problem. Here, we use $v_\epsilon[\cdot]$ to show the value function associated with ϵ . As we know for $\epsilon = 0$, the value of the game is infinity:

$$v_0[X] = \infty.$$

However, even when we choose a very small ϵ , apply the transformation given in Theorem 4, and use the g function above, the value of the game drops to a very small amount. Indeed, at just $\epsilon = 0.01$ (which is obtained at $\alpha = 0.056$), the value drops to

$$v_{0.01}[X] = 12.4.$$

Thus, we conclude that it is most likely not reasonable to pay more than 13 units of utility for a one-time shot at this gamble.

Algorithm 1 Calculation of $v[X] = E_Q[X]$ **INPUT:**

- 1) Distribution of X :
 - $X = \{x_1, x_2, \dots, x_r\}: x_1 < x_2 < \dots < x_r$
 - $P = \{p_1, p_2, \dots, p_r\}$
- 2) α

OUTPUT:

- 1) $v_\epsilon[X]$
- 2) ϵ

```

for  $k = 1 : r - 1$  do
   $q_k = g(\sum_{i=k}^r p_i) - g(\sum_{i=k+1}^r p_i)$ 
end for
 $q_r = g(p_r)$ 
 $v_\epsilon[X] = \sum_{k=1}^r x_k q_k$ 
 $\epsilon = \frac{1}{2} \sum_{i=1}^r |p_i - q_i|$ 

```

Next, let us look at the legal example provided in Section II. Remember that expected utility provided the following result:

$$E[X] = 10, \\ E[Y] = 0.35 \times 40 - 0.65 \times 5 = 10.75,$$

which suggests option Y (going to the court) is preferable. Now, by applying the transformation given in Theorem 4 and using the g function above, we notice that at

$$\epsilon^* = 0.0165,$$

which is obtained at $\alpha = 0.135$, we obtain

$$v[X] = v[Y] = 10.$$

That is, for any $\epsilon > 0.0165$, option X , i.e., accepting the settlement is preferable. Since $\epsilon^* = 0.0165$ is very small, we conclude that accepting the settlement, is most likely a preferable choice.

The above examples were for scenarios where we obtained a result other than what is proposed by expected utility. Nevertheless, expected utility provides reliable answers for many problems (even in non-LLN regimes). Indeed, for many such problems, we note that there is no ϵ for which the preferences change. This means that the above method produces the same result as the one obtained by the expected utility theory.

In other scenarios, the method produces a large ϵ^* , which again indicates agreement with expected utility theory. For example, suppose that an agent is choosing between winning $X = 95$ units of utility for sure and a gamble where she wins $Y = 100$ utility units with probability 80% and nothing ($Y = 0$) with probability 20%. In such a case, expected utility theory ($\epsilon = 0$) prefers X . Applying the above method, we obtain $\epsilon^* = 0.15$, which is a very large value, indicating that the result obtained by expected utility theory is reliable.

C. AN EXAMPLE OF A REAL-WORLD APPLICATION OF THE METHOD

Here, we would like to apply the systematic approach to the problem of angel and venture capital (VC) investment and compare the results with those obtained by a few other approaches. Suppose that an angel or a venture capital investment fund is being created to invest in technology startups. A fundamental question is how many startups the total available funds should be divided in? We will apply the proposed systematic change-of-measure-based approach to answer the question and compare the result with those from a few other approaches.

To formulate the problem, let L be the total number of companies in which the fund invests. Let $X^{(j)}, j = 1, 2, \dots, L$ be the total profit from the investment in the j th company assuming one unit of money being invested. For example, if the j^{th} company fails, we let $X^{(j)} = -1$. On the other hand, if the investor triples the invested amount, we let $X^{(j)} = 2$. For simplicity, we assume that the fund invests equal amounts in each company.

The first question that needs to be addressed is what the distribution of the $X^{(j)}$'s is. There are many works on the topic, for example, [71]–[74], and we adopt a model based on [71]. Specifically, we assume that the distribution of $X^{(j)}$ is as shown in Figure 9.




FIGURE 9. Distribution of U.S. venture returns between 2004 and 2013, adapted from [71].

Assume that the investor requires a minimum profit of 140% over the length of the investment, which is usually a few years for each startup. This seems to be consistent with the goals set by VCs and angel investors. Let us now try to address the question of how many companies the fund should invest in.

Expected Utility Approach: If we compute the expected utility, we get

$$E[X^{(j)}] = 1.52,$$

which implies that on average, the investor earns a profit of 150% dollars for each dollar she invests in a single company. This means that if we just want to use expected utility, even

a single startup suffices as $E[X^{(j)}] - 1.4 > 0$. Needless to say, VCs and angel investors are completely aware that the expected utility does not suffice for such risky investments. It is well known that a diversified portfolio is much less risky.

Limiting (Bounding) Loss Probability: One common approach is to require that the probability of a loss at the end of investment be lower than a given threshold, say 10%. More specifically,

$$P\left(\sum_{j=1}^L X^{(j)} < 0\right) < 0.1.$$

Assuming the investments are independent, we can calculate the required L using the distribution of the sum. The smallest L for which the probability of loss is smaller than 10% is

$$L = 12.$$

Thus, the investor could choose $L = 12$. It is worth noting that due to the discreteness of the distribution, the loss probability is not monotonic for a small L . For example, the loss probability is larger than 10% for $L = 13$, but it becomes smaller than 10% for all $L \geq 14$. Thus, the investor may choose $L \geq 14$ to be on the safe side.

Proposed Systematic Approach Based on Consistent Functions: We can follow the systematic approach discussed in the previous section. Specifically, we look at

$$v_\epsilon \left[\frac{1}{L} \sum_{j=1}^L X^{(j)} - 1.4 \right],$$

and for each L obtain the range of values of ϵ where the above quantity is positive. Remember that this range gives us stability, and we want it to be larger. If we require a 0.05 tolerance (i.e., $\epsilon^* = 0.05$), we obtain

$$L \geq 12.$$

Therefore, the proposed systematic approach provides similar results to the approach based on limiting loss probability in this case. However, the systematic approach based on the change-of-probability measures has some desirable properties:

- First, it is a general approach that can be applied to any situation, i.e., in LLN and non-LLN regimes.
- Second, it is a systematic approach that can be easily programmed into AI decision-making.
- Third, the proposed approach is based on the entire probability distribution, while approaches such as limiting the loss probability only look at part of the distribution (e.g., the part resulting in a loss).
- Finally, as discussed before, combining different approaches seems to be a reasonable approach to probabilistic decision-making, and the proposed change-of-measure approach can be a key component in that regard.

VI. DISCUSSION AND FURTHER RESEARCH

Here, we provide some discussions and comments on a few potential avenues for further research. Decision-making under non-LLN regimes is an important area to investigate, and its importance is growing with the rise of AI. Almost all high-impact decisions are in this category. Unlike LLN regimes where expected utility theory provides a relatively satisfying answer, the problem of decision-making under non-LLN regimes is multidimensional and complex, and each specific instance might require specific consideration. Therefore, it is very unlikely that a single approach can provide all the answers. With those considerations in mind, this article aimed to provide a framework based on the change-of-probability measures to shed light on some aspects of this important area. We observed that the proposed method provides satisfactory results for some problems and was able to provide some insights.

We believe the presented material potentially provides several avenues for further research. First, the proposed approach should be applied to different problems within different contexts. Undoubtedly, such efforts will reveal shortcomings, and this could help bring us closer to more comprehensive approaches. Different specific non-LLN problems in engineering, computer science, philosophy, economics, and other social sciences can be considered to test and improve the proposed approach. This could help yield more insights into the types of problems for which different approaches are more effective.

It is important to note that the proposed method addresses a single issue: the non-LLN nature of some decision-making problems. For more comprehensive decision-making, a very promising lead could be to combine the proposed methodology with other techniques, such as those developed for risk management.

Next, there is much flexibility within the proposed framework. The approach used in Sections V-A, V-B, and V-C is only one of the possible approaches under the proposed framework. Different approaches using different forms of change-of-measure operations might be more effective for some problems.

There are many problems that can be considered from a mathematical perspective. Here, we mention a few: First, it is possible to add to the required properties to make the set of change-of-measures more restrictive. Second, the stated results can be extended and improved. More properties of ϵ -consistent change-of-measure policies could be proved.

Finally, in the algorithms that we have provided, we focused on the cases in which there is only one change of priorities for $\epsilon^* < 0.05$. However, it is possible that even the second option is not stable enough, such that the preferred option changes very quickly by increasing ϵ slightly from the ϵ^* . For these cases, more sophisticated strategies for optimal decision-making must be developed, which is an interesting question to investigate in further research.

VII. CONCLUSION

In this article, we considered the fundamental problem of decision-making under non-LLN regimes. We first introduced BE as a principled way to address the accepted principle of ignoring negligible probabilities. BE provided satisfactory answers and insights regarding some aspects of decision-making under non-LLN regimes. Pointing out some shortcomings of BE, we then extended the approach to a much more general framework of change-of-probability measures. The proposed theory can be considered to be a generalization of expected utility theory in two directions. First, it was shown that as the number of repetitions increases, the results derived from the proposed theory converges to those from expected utility theory. Second, when the distortion parameter, ϵ , is zero, the proposed theory again becomes identical to expected utility. Finally, we suggested a systematic approach to applying the theory and showed that it produces satisfactory results for some examples. The proposed paradigm can be applied to different high-impact and non-LLN decisions made by humans or AI in business, economics, medicine, and computer science, to name a few. However, due to the complexity and multidimensionality of such problems, there may be limitations in the proposed method that need to be carefully investigated. Hence, we noted that this article could potentially lead to considerable further research.

ACKNOWLEDGMENT

The authors would like to thank Ali Khezeli (Tarbiat Modares University) for very helpful feedback and comments on the paper.

REFERENCES

- [1] M. Peterson, *An Introduction to Decision Theory*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2017.
- [2] N. D. Nguyen, T. Nguyen, and S. Nahavandi, "System design perspective for human-level agents using deep reinforcement learning: A survey," *IEEE Access*, vol. 5, pp. 27091–27102, Nov. 2017.
- [3] Z. Xu, L. Cao, and X. Chen, "Learning to learn: Hierarchical meta-critic networks," *IEEE Access*, vol. 7, pp. 57069–57077, May 2019.
- [4] G. Sun, Z. T. Gebrekidan, G. O. Boateng, D. Ayepah-Mensah, and W. Jiang, "Dynamic reservation and deep reinforcement learning based autonomous resource slicing for virtualized radio access networks," *IEEE Access*, vol. 7, pp. 45758–45772, Apr. 2019.
- [5] C. Wu, B. Ju, Y. Wu, X. Lin, N. Xiong, G. Xu, H. Li, and X. Liang, "UAV autonomous target search based on deep reinforcement learning in complex disaster scene," *IEEE Access*, vol. 7, pp. 117227–117245, Aug. 2019.
- [6] N. C. Luong, D. T. Hoang, S. Gong, D. Niyato, P. Wang, Y.-C. Liang, and D. I. Kim, "Applications of deep reinforcement learning in communications and networking: A survey," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 4, pp. 3133–3174, 4th Quart., 2019.
- [7] G. Li, R. Gomez, K. Nakamura, and B. He, "Human-centered reinforcement learning: A survey," *IEEE Trans. Human-Mach. Syst.*, vol. 49, no. 4, pp. 337–349, Aug. 2019.
- [8] Q. Mao, F. Hu, and Q. Hao, "Deep learning for intelligent wireless networks: A comprehensive survey," *IEEE Commun. Surveys Tuts.*, vol. 20, no. 4, pp. 2595–2621, 4th Quart., 2018.
- [9] M. Mahmud, M. S. Kaiser, A. Hussain, and S. Vassanelli, "Applications of deep learning and reinforcement learning to biological data," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 29, no. 6, pp. 2063–2079, Jun. 2018.
- [10] A. Tversky, "A critique of expected utility theory: Descriptive and normative considerations," *Erkenntnis*, vol. 9, no. 2, pp. 163–173, Jun. 1975.
- [11] G. Wu, "The strengths and limitations of expected utility theory," *Med. Decis. Making*, vol. 16, no. 1, pp. 9–10, Feb. 1996.
- [12] P. J. Schoemaker, "The expected utility model: Its variants, purposes, evidence and limitations," *J. Econ. Literature*, pp. 529–563, Jun. 1982.
- [13] N. N. Taleb and R. Douady, "Mathematical definition, mapping, and detection of (anti)fragility," *Quant. Finance*, vol. 13, no. 11, pp. 1677–1689, Nov. 2013.
- [14] N. N. Taleb, *The Black Swan: The Impact Of The Highly Improbable*, vol. 2. New York, NY, USA: Random House, 2007.
- [15] M. Frize, L. Yang, R. C. Walker, and A. M. O'Connor, "Conceptual framework of knowledge management for ethical decision-making support in neonatal intensive care," *IEEE Trans. Inf. Technol. Biomed.*, vol. 9, no. 2, pp. 205–215, Jun. 2005.
- [16] X. Fu, "Application of artificial intelligence technology in medical cell biology," in *Proc. Int. Conf. Robots Intell. Syst. (ICRIS)*, Jun. 2019, pp. 401–404.
- [17] M. A. Salichs and M. Malfaz, "A new approach to modeling emotions and their use on a decision-making system for artificial agents," *IEEE Trans. Affect. Comput.*, vol. 3, no. 1, pp. 56–68, Jan. 2012.
- [18] M. Frutos-Pascual and B. G. Zapirain, "Review of the use of AI techniques in serious games: Decision making and machine learning," *IEEE Trans. Comput. Intell. AI Games*, vol. 9, no. 2, pp. 133–152, Jun. 2017.
- [19] J. Lu, Z. Yan, J. Han, and G. Zhang, "Data-driven decision-making (D³M): Framework, methodology, and directions," *IEEE Trans. Emerg. Topics Comput. Intell.*, vol. 3, no. 4, pp. 286–296, Aug. 2019.
- [20] D. Chumachenko, I. Meniaikov, K. Bazilevych, and T. Chumachenko, "On intelligent decision making in multiagent systems in conditions of uncertainty," in *Proc. 11th Int. Sci. Practical Conf. Electron. Inf. Technol. (ELIT)*, Sep. 2019, pp. 150–153.
- [21] R. R. Yager, "Modeling prioritized multicriteria decision making," *IEEE Trans. Syst., Man, Cybern. B, Cybern.*, vol. 34, no. 6, pp. 2396–2404, Dec. 2004.
- [22] K. Huang, C. Zhou, Y. Qin, and W. Tu, "A game-theoretic approach to cross-layer security decision-making in industrial cyber-physical systems," *IEEE Trans. Ind. Electron.*, vol. 67, no. 3, pp. 2371–2379, Mar. 2020.
- [23] Y. Li, W. Han, and Y. Wang, "Deep reinforcement learning with application to air confrontation intelligent decision-making of manned/unmanned aerial vehicle cooperative system," *IEEE Access*, vol. 8, pp. 67887–67898, Apr. 2020.
- [24] Y. A. Basallo, V. E. Senti, and N. M. Sanchez, "Artificial intelligence techniques for information security risk assessment," *IEEE Latin Amer. Trans.*, vol. 16, no. 3, pp. 897–901, Mar. 2018.
- [25] G. Kłosowski and A. Gola, "Risk-based estimation of manufacturing order costs with artificial intelligence," in *Proc. Federated Conf. Comput. Sci. Inf. Syst. (FedCSIS)*, Oct. 2016, pp. 729–732.
- [26] D. Kahneman and A. Tversky, "Prospect theory: An analysis of decision under risk," *Econometrica*, vol. 47, no. 2, pp. 263–292, Mar. 1979.
- [27] A. Tversky and D. Kahneman, "Advances in prospect theory: Cumulative representation of uncertainty," *J. Risk Uncertainty*, vol. 5, no. 4, pp. 297–323, Oct. 1992.
- [28] K. D. Edwards, "Prospect theory: A literature review," *Int. Rev. Financial Anal.*, vol. 5, no. 1, pp. 19–38, Jan. 1996.
- [29] P. C. Fishburn, "Subjective expected utility: A review of normative theories," *Theory Decis.*, vol. 13, no. 2, pp. 139–199, Jun. 1981.
- [30] D. E. Bell, H. Raiffa, and A. Tversky, *Decision Making: Descriptive, Normative, and Prescriptive Interactions*. Cambridge, U.K.: Cambridge Univ. Press, 1988.
- [31] M. Roubens and J. Kacprzyk, *Non-Conventional Preference Relations in Decision Making*. Berlin, Germany: Springer-Verlag, 1988.
- [32] R. Briggs. (2014). *Normative Theories of Rational Choice: Expected Utility*. [Online]. Available: <https://plato.stanford.edu/entries/rationality-normative-utility/>
- [33] L. Buchak and L. M. Buchak, *Risk and Rationality*. London, U.K.: Oxford Univ. Press, 2013.
- [34] J. Quiggin, *Generalized Expected Utility Theory: The Rank-Dependent Model*. Dordrecht, The Netherlands: Springer, 1993.
- [35] M. Colyvan, "Relative expectation theory," *J. Philos.*, vol. 105, no. 1, pp. 37–44, 2008.
- [36] D. Schmeidler, "Subjective probability and expected utility without additivity," *Econometrica J. Econ. Soc.*, vol. 57, no. 3, pp. 571–587, May 1989.

[37] M. S. Pfaff, J. L. Drury, G. L. Klein, L. More, S. P. Moon, and Y. Liu, "Weighing decisions: Aiding emergency response decision making via option awareness," in *Proc. IEEE Int. Conf. Technol. Homeland Secur. (HST)*, Nov. 2010, pp. 251–257.

[38] M. Cashmore, M. Fox, D. Long, D. Magazzeni, and B. Ridder, "Opportunistic planning in autonomous underwater missions," *IEEE Trans. Autom. Sci. Eng.*, vol. 15, no. 2, pp. 519–530, Apr. 2018.

[39] O. Efthimiou, "Practical guide to the meta-analysis of rare events," *Evidence Based Mental Health*, vol. 21, no. 2, pp. 72–76, May 2018.

[40] P. Goodwin and G. Wright, "The limits of forecasting methods in anticipating rare events," *Technological Forecasting Social Change*, vol. 77, no. 3, pp. 355–368, Mar. 2010.

[41] C. Arney, K. Coronges, H. Fletcher, J. Hagen, K. Hutchinson, A. Moss, and C. Thomas, "Using rare event modeling & networking to build scenarios and forecast the future," in *Proc. IEEE 2nd Netw. Sci. Workshop (NSW)*, Apr. 2013, pp. 31–36.

[42] I. B. Sperstad and E. S. Kiel, "Development of a qualitative framework for analysing high-impact low-probability events in power systems," in *Proc. Saf. Rel.-Safe Societies Changing World (ESREL)*, Trondheim, Norway, 2018, pp. 1599–1608.

[43] K. J. Hole, "Management of hidden risks," *Computer*, vol. 46, no. 1, pp. 65–70, Jan. 2013.

[44] K. J. Hole and L.-H. Netland, "Toward risk assessment of large-impact and rare events," *IEEE Secur. Privacy Mag.*, vol. 8, no. 3, pp. 21–27, May/Jun. 2010.

[45] P. Chalasani and S. Jha, "Steven Shreve: Stochastic calculus and finance," *Lect. Notes*, Oct. 1997. [Online]. Available: <http://blog.yg.pe.kr/attachment/ek3.pdf>

[46] T. H. Chan, S. Hranilovic, and F. R. Kschischang, "Capacity-achieving probability measure for conditionally Gaussian channels with bounded inputs," *IEEE Trans. Inf. Theory*, vol. 51, no. 6, pp. 2073–2088, Jun. 2005.

[47] K. Todros and A. O. Hero, "Robust multiple signal classification via probability measure transformation," *IEEE Trans. Signal Process.*, vol. 63, no. 5, pp. 1156–1170, Mar. 2015.

[48] L. Jin, R. Mesiar, and R. R. Yager, "Melting probability measure with OWA operator to generate fuzzy measure: The crescent method," *IEEE Trans. Fuzzy Syst.*, vol. 27, no. 6, pp. 1309–1316, Jun. 2019.

[49] M. Peterson, "The St. Petersburg paradox," in *The Stanford Encyclopedia Philosophy*, E. N. Zalta, Ed. Stanford Univ. Press, Fall 2019.

[50] D. Bernoulli, "Exposition of a new theory on the measurement of risk," , vol. 12, no. 1, pp. 23–36, Jan. 1954.

[51] K. Menger, "Das unsicherheitsmoment in der wertlehre: Betrachtungen im angeschliss an das sogenannte petersburger spiel," *Zeitschrift für Nationalökonomie/J. Econ.*, vol. 5, no. 4, pp. 459–485, Aug. 1934.

[52] I. Hacking, "Strange expectations," *Philos. Sci.*, vol. 47, no. 4, pp. 562–567, Dec. 1980.

[53] J. D. Hey, T. M. Neugebauer, and C. M. Pasca, "Georges-Louis Leclerc de Buffon's 'Essays on moral arithmetic,'" in *The Seltzer School of Behavioral Economics*. New York, NY, USA: Springer-Verlag, 2010, pp. 245–282.

[54] T. M. Cover, "On the St. Petersburg paradox," in *Proc. IEEE Int. Symp. Inf. Theory*, Jul. 2011, pp. 1758–1761.

[55] C. Seidl, "The St. Petersburg paradox at 300," *J. Risk Uncertainty*, vol. 46, no. 3, pp. 247–264, Jun. 2013.

[56] N. J. J. Smith, "Is evaluative compositionality a requirement of rationality?" *Mind*, vol. 123, no. 490, pp. 457–502, Aug. 2014.

[57] O. Peters, "The time resolution of the st petersburg paradox," *Phil. Trans. Roy. Soc. A: Math., Phys. Eng. Sci.*, vol. 369, no. 1956, pp. 4913–4931, Dec. 2011.

[58] P. Billingsley, *Probability and Measure*, 3rd ed. Hoboken, NJ, USA: Wiley, 1986.

[59] W. Feller, *An Introduction to Probability Theory and Its Applications*. Hoboken, NJ, USA: Wiley, 1957.

[60] A. Pine, B. Seymour, J. P. Roiser, P. Bossaerts, K. J. Friston, H. V. Curran, and R. J. Dolan, "Encoding of marginal utility across time in the human brain," *J. Neurosci.*, vol. 29, no. 30, pp. 9575–9581, Jul. 2009.

[61] J. Mumpower, "An analysis of the de minimis strategy for risk management," *Risk Anal.*, vol. 6, no. 4, pp. 437–446, Dec. 1986.

[62] M. D. Adler, "Why de minimis?" *Inst. Law Econ.*, Univ. Pennsylvania, Philadelphia, PA, USA, Res. Paper 07–12, Jun. 2007, pp. 7–26.

[63] M. Peterson, "What is a de minimis risk?" *Risk Manage.*, vol. 4, no. 2, pp. 47–55, Apr. 2002.

[64] T. Neugebauer, "Moral impossibility in the petersburg paradox: A literature survey and experimental evidence," *Luxembourg School Finance Res. Work. Paper Ser.*, vol. 10, no. 174, pp. 1–43, 2010.

[65] A. Hájek, "Unexpected expectations," *Mind*, vol. 123, no. 490, pp. 533–567, Apr. 2014.

[66] J. W. Baker, M. Schubert, and M. H. Faber, "On the assessment of robustness," *Struct. Saf.*, vol. 30, no. 3, pp. 253–267, May 2008.

[67] J. Ruhl, "Managing systemic risk in legal systems," *Ind. LJ*, vol. 89, no. 2, p. 559, Jun. 2014.

[68] S. Martínez-Jaramillo, O. P. Pérez, F. A. Embriz, and F. L. G. Dey, "Systemic risk, financial contagion and financial fragility," *J. Econ. Dyn. Control*, vol. 34, no. 11, pp. 2358–2374, Nov. 2010.

[69] M. Mishra, D. Sidoti, G. V. Avvari, P. Mannaru, D. F. M. Ayala, K. R. Patti-pati, and D. L. Kleinman, "A context-driven framework for proactive decision support with applications," *IEEE Access*, vol. 5, pp. 12475–12495, May 2017.

[70] J. S. Rosenthal, *A First Look At Rigorous Probability Theory*. Singapore: World Scientific, 2006.

[71] H. Swildens and E. Yee. (Feb. 7, 2017). *The Venture Capital Risk and Return Matrix*. [Online]. Available: <http://www.industryventures.com/the-venture-capital-risk-and-return-matrix/>

[72] The Quantified VC. (Aug. 23, 2018). *How to Win in Venture Capital: Focus on the Fat Tails*. [Online]. Available: <https://blog.usejournal.com/power-laws-in-venture-capital-why-the-long-tail-matters-22e057c6fa34>

[73] Guillem. (May 18, 2017). *Understanding the Nature of Venture Capital Returns*. [Online]. Available: <https://medium.com/guillem-sague/understanding-the-nature-of-venture-capital-returns-95846c65c049>

[74] R. Wiltbank. (Oct. 13, 2012). *Angel Investors do Make Money, Data Shows 2.5x Returns Overall*. [Online]. Available: <https://techcrunch.com/2012/10/13/angel-investors-make-2-5x-returns-overall/>



SAEDEE ENAYATI received the B.Sc. degree from Shahid Beheshti University, in 2012, and the M.Sc. degree from Tarbiat Modares University, Tehran, Iran, in 2015. She is currently pursuing the Ph.D. degree with the Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, MA, USA. She was a Research Assistant, from 2016 to 2019. Her research interests include stochastic processes and wireless network analysis.



HOSSEIN PISHRO-NIK (Member, IEEE) received the B.Sc. degree from the Sharif University of Technology, Tehran, Iran, and the M.Sc. and Ph.D. degrees from the Georgia Institute of Technology, all in electrical and computer engineering. He is currently a Professor with the Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, MA, USA. His research interests include information-theoretic privacy and security, error control coding, vehicular/UAV networks, and mathematical analysis of wireless networks. He was a recipient of the NSF Faculty Early Career Development (CAREER) Award, the Outstanding Junior Faculty Award from UMass, and the Outstanding Graduate Research Award from the Georgia Institute of Technology. He is an Associate Editor of the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS.