Centro UC
CLAPES UC
Centro Latinoamericano de Políticas Económicas y Sociales

A CRITICAL ASSESSMENT OF THE KNIGHT-MARKOWITZ CONCEPTION OF RISK

Septiembre, 2026

AUTORES
Arturo Cifuentes | Investigador Principal CLAPES UC
Ventura Charlin | Investigadora Asociada CLAPES UC

DOC.DE TRABAJO
Nº171

A CRITICAL ASSESSMENT OF THE KNIGHT-MARKOWITZ CONCEPTION OF RISK

ABSTRACT

This paper examines Knight's distinction between risk and uncertainty, introduced in 1921. We show that prior to Knight classical economists employed risk and uncertainty in accordance with common usage: risk referred to the possibility of adverse outcomes. Many technical and scientific disciplines still employ these terms this way. Knight's redefinition of risk, as randomness with known probabilities, however, combined favorable and unfavorable outcomes into one concept—a misconception that, on its own, remained analytically inert for three decades, since Knight supplied no formula to operationalize it. The error became consequential only when Markowitz (1952), provided Knight's concept with the computational tool it had lacked—thus turning a conceptual mistake into a powerful tool. Financial engineering, a younger discipline, has instead returned to pre-Knightian conceptions of risk, adopting metrics focused exclusively on unfavorable outcomes. Economics would benefit from realigning with this approach, correcting a decision that yielded neither theoretical insights nor practical results.

Keywords: Risk; Uncertainty; Knight; Markowitz

JEL Codes: A12; B41; B26; C18; D81

Introduction

Risk is part of life, and most of the decisions we make implicitly involve it in varying degrees. Not surprisingly, people have a clear and intuitive sense of what this concept means. By risk, we refer to the possibility of experiencing harm or danger —that is, the possibility of an adverse outcome. Uncertainty, while related, is a different concept: it refers to our ignorance —total or partial— regarding a future outcome, regardless of whether that outcome is favorable or not. The present paper is mainly concerned with risk, and with the consequences of how economics chose to redefine it and measure it.

A headline from *The New York Times* reflects the common interpretation of risk clearly: '8 Factors That Can Raise Your Heart Disease Risk,' from the November 7, 2024 edition. Risk, in this and virtually every other context outside economics, refers unambiguously to the possibility of something undesirable occurring.

More formal definitions validate this interpretation. The Merriam-Webster dictionary (online version, consulted May 5, 2025) defines risk as 'possibility of loss or injury.' Other dictionaries, such as the American Heritage, offer nearly identical definitions.

Scientific and technical disciplines employ the term in a manner fully consistent with common usage. Examples abound across psychology (Tversky and Kahneman, 1974), civil engineering (Faber and Stewart, 2003), mechanical engineering (Ayyub, 2015), meteorology (Wilks and Horowitz, 2014), sports medicine (Bahr and Holme, 2003), earthquake engineering (Ellingwood and Kinali, 2009), insurance (Botticini et al., 2023), chemistry (Duke and Taggart, 2000), biology (Barratt et al., 2023), applied statistics (Albarrán et al., 2025), decision sciences (Cox, 2012), and the pharmaceutical industry (Scherer, 2007). In all these cases, risk refers to potentially negative or undesirable situations —property damage caused by hurricanes, injuries from snowboarding accidents, impairment of critical infrastructure due to earthquakes, impacts on human health from water pollution, and so on.

Philosophical discussions of risk also use the term in line with common parlance. Hansson (2004), in *Philosophical Perspectives on Risk*, opens by defining risk as associated with unwanted events, offering the example that ‘lung cancer is one of the major risks that affect smokers.’ The Stanford Encyclopedia of Philosophy concurs, describing risk as ‘an unwanted event that may or may not occur’ (Hansson, 2023).

The picture that emerges from this survey is remarkably consistent: across everyday language, many popular dictionaries, scientific and technical disciplines, and philosophical inquiry, risk refers to the possibility of adverse outcomes. Economics, however, stands apart. Why did it adopt a definition of risk that departs not only from common understanding but from the usage of every other discipline that employs the term? The question is far from trivial. Hausman (2023), for example, has cautioned that while economists are entitled to assign technical meanings to concepts drawn from ordinary language, that entitlement is conditional: such departures carry costs as well as potential benefits —it is a decision with consequences.

The present paper examines the origins of the particular definition of risk adopted by economics —one that continues to prevail in current economic thinking— and argues that its costs have substantially outweighed its benefits. We make the case that this definition, introduced by Knight, has had lasting negative effects on the subsequent development of economics. The argument identifies two distinct and sequential problems. The first is definitional: Knight's redefinition of risk as randomness with known probabilities departed from common usage and from the conventions of every other discipline that employs the term. Specifically, it collapses favorable and unfavorable outcomes into a single concept, without any compelling justification. The second is a problem of measurement: when Markowitz (1952) eventually operationalized Knight's definition by adopting the standard deviation of returns as its quantitative expression, he introduced a metric whose symmetric treatment of favorable and unfavorable outcomes was inconsistent with any practically meaningful conception of risk —though, notably, entirely consistent with Knight's own definition. In this sense, rather than correcting the conceptual problem, Markowitz's formula gave Knight's flawed definition the tool it had lacked, and in doing so helped entrench the misunderstanding. These two problems —a flawed definition, and the adoption of a metric that operationalized it— are analytically distinct, and keeping them separate is essential to understanding both how the error arose and what would be required to correct it.

The Knightian Turn

Frank Knight (1885–1972) is considered a towering figure among economists although his name, unlike some of his peers, is largely unknown outside the discipline. His reputation was established in 1921 with the publication of his magnum opus *Risk, Uncertainty and Profit* (Knight, 1921). The book was well received and has been republished several times; it is now widely available in physical and electronic form. Its goal was to elaborate on the reasons an entrepreneur can make a profit in a capitalist system. Our concern here, however, is a different point: the distinction Knight drew between risk and uncertainty. He claimed that there was confusion regarding the meaning of these two terms; more precisely, he stated that these two terms 'have never been properly separated.' Without providing any evidence to support this claim, he took it upon himself to rectify this alleged confusion.

Knight adopted the term *risk* for situations of randomness in which the probabilities of the potential outcomes or possible scenarios are known. Rolling a fair die is the classical example: the probability of any given face appearing is one-sixth. He reserved the term *uncertainty* for situations of randomness in which those probabilities are unknown, for instance, rolling a die known to be loaded, but without any knowledge of how it was manipulated.

One key feature of Knight's definition of risk is that it does not distinguish between positive and negative outcomes, that is, between favorable and unfavorable results.

Notably, Knight neither proposed nor even suggested a formula for calculating risk according to his definition. His book is essentially a long essay written in prose, without equations or formulas, though it does include a handful of simple graphs. This is perhaps unsurprising, and consistent with the view held by some scholars that for Knight, risk and uncertainty were matters of judgment rather than precise calculation (see, for example, Cowan, 2016). Judgment, by its nature, is difficult to articulate through a mathematical expression. This absence of any quantitative instrument left Knight's definition of risk without an operational component. In essence, the concept remained analytically inert until Markowitz (1952) supplied the missing formula three decades later —a formula that, as the following sections will show, introduced a second and independent problem of its own. In supplying that formula, Markowitz did more than simply operationalize an inert idea: he gave Knight's flawed definition a mathematical vehicle for propagation, one that a merely discursive definition could never have achieved. It was this pairing of definition and formula, rather than either element on its own, that allowed a definition arguably at odds with common usage to become entrenched as the standard conception of risk within economics.

In fairness to Knight, he did identify a genuinely important issue: the distinction between what he called measurable and non-measurable probabilities, or more formally, the need to recognize the uncertainty inherent in some probability estimates attached to potential scenarios. Today, this issue can be addressed with relative ease using techniques such as Monte Carlo simulations combined with sensitivity analyses, methodologies that were unavailable in Knight's time. It is understandable, then, that the problem of dealing with probabilities that could not be determined with certainty —and had instead to be estimated— might have seemed insurmountable in his era. Unfortunately, his discussion of it became obscured by the broader ambiguities he introduced.

Two further points are worth noting, though both are largely tangential to the discussion that follows. First, Knight claimed, again without much supporting evidence, that people ordinarily reserved the term *uncertainty* for favorable outcomes only. Common usage does not appear to support this view. Second, and somewhat ironically, Knight himself frequently used the terms risk and uncertainty throughout his book in their conventional sense, rather than according to his own definitions. This is precisely the inconsistency Hausman (2023) warns against: the license to redefine a term carries the obligation to use the redefined concept consistently —one cannot proscribe the everyday meaning and then revert to it when it is convenient. (NOTE: the Appendix provides a useful timeline with the evolution of the concept of risk.)

Was Knight's Redefinition of Risk Justified?

When Knight asserted that there was confusion and ambiguity regarding the meaning of risk and uncertainty in common usage, and thus saw the need to articulate precise definitions of these terms, he presented no proof to support this opinion. One might interpret this absence of evidence not as a shortcoming of his exposition, but simply as a sign that the confusion was so obvious that it required no illustration. The reality, however, points in a different direction: there was simply no evidence to support Knight's view. In fact, many economists long before Knight had been using risk and uncertainty in ways fully consistent with their common meaning. Adam Smith (1723–1790) is perhaps the clearest example.

In *The Wealth of Nations* (Smith, 1776), Smith used the term *risk* numerous times, always in reference to potentially undesirable outcomes:

> 'risk of population falling below the level of subsistence…'
>
> 'risk of a suspension of cash payments…'
>
> 'risk of falling a victim to the vindictive spirit…'

In other passages, he invoked the term in relation to the possibility of tax increases, potential losses due to fire or capture during sea transport, and the dangers of keeping cash at home. Moreover, on multiple occasions he paired the word *risk* with the word *trouble* in the same sentence, further underscoring the harmful nature of the situations he was describing. In every case, his usage was entirely consistent with the common interpretation of the term.

Smith also employed *uncertainty* in a way that conveyed lack of clarity or ignorance regarding a future event:

> ‘the uncertainty of recovering his money…’
>
> ‘the measure of value comes to be liable to the same sort of uncertainty to which all other weights and measures are commonly exposed…’
>
> ‘the uncertainty of getting those bank notes exchanged for gold…’

Other economists tell a similar story. Richard Cantillon (1680–1734), in his Essay on the *Nature of Trade in General* (Cantillon, 1755), wrote of ‘the risk of not receiving his rent’ and ‘the risk of losing their capital,’ alongside ‘uncertainty since their customers may leave them from one day to the next’ and ‘uncertainty with regard to their earnings.’ David Ricardo (1772–1823), in his *Principles of Political Economy and Taxation* (Ricardo, 1817), referred to ‘the risk of ruin.’ And John Stuart Mill (1806–1873), in his Principles of Political Economy (Mill, 1870), recognized ‘the uncertainty of all things future.’

In short, notwithstanding Knight's assertion that confusion existed regarding the meaning of risk and uncertainty, there appears to be no evidence to validate this view. On the contrary, there is ample record that highly regarded economic treatises used both terms in full agreement with common parlance, and none of their authors was ever accused of creating confusion by relying on language whose meaning was unclear.

It is also worth noting that John Maynard Keynes (1883–1946), a contemporary of Knight, apparently saw no such confusion. Unlike Knight, Keynes moved directly to a more substantive issue that had somehow escaped Knight's attention: the fact that in some situations, the potential outcomes themselves —not merely their probabilities— are unknown. This point has been taken up more recently by several authors (e.g., Packard et al., 2021; Bhide, 2021), who refer to Knight's version of uncertainty as *aleatory* and reserve the term *epistemic uncertainty* for the situations Keynes had in mind. Notably, neither Knight nor Keynes used these terms to describe the types of uncertainty they were actually discussing.

Some Post-Knight Considerations

Apparently not through any effort initiated by Knight himself, his definition of uncertainty eventually became known as *Knightian uncertainty*. The origins of this label merit some attention, as does the broader reception of Knight's ideas in the decades following the publication of his book.

Knight's book was reviewed in 1922, a year after its publication, by Wesley Mitchell (1922) and G. P. Watkins (1922) in two prestigious economics journals. Neither review was particularly enthusiastic, though Watkins engaged more substantively with the risk/uncertainty distinction, even if his treatment was critical rather than approving. Regardless, the book's reputation received a significant boost following its republication in 1933, an outcome that resulted directly from Lionel Robbins having assigned it as a core text in one of his courses at the London School of Economics.

Oliver G. Wood Jr.'s *Evolution of the Concept of Risk*, published in the *Journal of Risk and Insurance* in 1964 (Wood, 1964), appears to be the first article devoted explicitly to tracing the historical development of the concept of risk. From the perspective of the present article, Wood's historical survey lends further support to the argument that risk, prior to Knight's intervention, was consistently understood as the possibility of experiencing adverse or undesirable outcomes, and that Knight's redefinition therefore represented a genuine departure from that tradition rather than a clarification of it. Tellingly, although Wood credits Knight with drawing a distinction between risk and uncertainty, he concludes that Knight's contribution ‘does not appear to support any of the definitions of risk given earlier in this paper.’ In short, Wood regarded Knight's framework as failing to provide a satisfactory definition of risk. This is perhaps unsurprising given that Wood ultimately recommends ‘chance of loss’ as the most suitable definition —a conclusion consistent with both the pre-Knightian tradition and the everyday usage that Knight himself had critiqued. Also worth noting is that the term ‘Knightian,’ and specifically ‘Knightian uncertainty,’ does not appear anywhere in Wood's article.

The earliest documented use of the term *Knightian uncertainty* in the published academic literature appears to be Schmalensee (1974), who employed it twice —on pages 269 and 270 of his paper— without explanation or definition, treating it as an expression already familiar to his readers. This casual usage strongly suggests that the term had been in circulation somewhat earlier, at least informally or in working papers, though no earlier published instance has yet been identified with certainty. It is worth noting that Schmalensee used the term in the context of consumer behavior in disequilibrium markets, which is quite different from the decision-theory contexts in which Bewley (2002) later employed it. Interestingly, LeRoy and Singell (1987), despite devoting their entire article to a close analysis of Knight's distinction between risk and uncertainty, never once use the word *Knightian*, let alone the compound term *Knightian uncertainty*.

Two earlier works addressed Knight's distinction between risk and uncertainty directly, though neither employed the term *Knightian* uncertainty: Marschak (1938) and Makower and Marschak (1938) explicitly addressed Knight's definitions, interpreting his notion of risk as the dispersion of the frequency distribution of alternative future outcomes, and his uncertainty as the degree of ignorance about that distribution. Notably, their engagement was not merely conceptual: they attempted to operationalize Knight's ideas within a mathematical framework, something Knight himself had failed to do. Indeed, Makower and Marschak (1938) hinted that risk was associated with the coefficient of variation —a well-known statistical metric that depends on the standard deviation— but did not propose a precise formula for computing it.

Later, Turvey, writing one of the reviews of *Expectation in Economics* (Shackle, 1949), in *Economica*, employed the terms Knight's risk and Knight's uncertainty (Turvey et al., 1949). He used the former when actuarial frequency distributions were available and could serve as a basis to establish definite numerical probabilities, and reserved the latter for cases in which no such data existed and probabilities had instead to be estimated subjectively.

The formal consolidation of the term *Knightian uncertainty* in the academic literature is generally associated with two works: Bewley's, 1986 Cowles Foundation working paper (later published in 2002) titled *Knightian Decision Theory*, and LeRoy and Singell's, 1987 paper in the Journal of Political Economy. Notably, while Bewley used the term Knightian explicitly, LeRoy and Singell did not —yet their rigorous analysis of Knight's framework contributed substantially to its entrenchment in the economics lexicon. Today it is in wide use: Google Scholar identifies more than 500 articles dealing with Knightian uncertainty since 2000 alone.

Associating Knight's name with this particular definition of uncertainty carries an implicit acknowledgement that deserves a closer examination. By assigning it a special label, the economics community was, perhaps unwittingly, conceding that the definition was in some sense peculiar, at odds with the conventional interpretation of the concept, and therefore in need of a distinguishing label. Knight himself, of course, saw no such need. He regarded his definition of uncertainty as clarifying a confusion he believed to be widespread. As argued at length in the preceding section, however, that confusion does not appear to have existed.

In any event, discussions of Knight's book remained exclusively conceptual in character. This is unsurprising given that Knight failed to propose or develop specific analytical tools, such as mathematical formulas, to operationalize his definitions of risk and uncertainty. As a consequence, no one attempted to perform calculations based on his framework, much less to compare the results with estimates derived from competing definitions.

Risk Between Knight and Markowitz: Influence Without Operationalization

Knight's redefinition of risk reshaped the vocabulary of economics, but it did so without reshaping its analytical practice. Knight himself provided no mathematical tool for quantifying risk under his definition; as a result, among academics his framework remained discursive —invoked, but never operationalized. It could not be applied to analyze data, used to build models, or tested against empirical implications.

Among financial practitioners and investment writers, however, a rather different picture emerged. The evidence from Hartzmark and Solomon (2026), who undertook a systematic close reading of academic finance articles, practitioner books, and financial newspapers from the early twentieth century, points to a striking conclusion: the dominant conception of risk among practitioners during this period was quite different from Knight's definition —and, in certain important respects, closer to the pre-Knightian understanding of risk as the possibility of adverse outcomes, even if expressed in imprecise and often confused terms.

The prevailing view among financial practitioners and investment writers during the, 1920s, 1930s, and, 1940s was that risk meant, essentially, the possibility of losing money over the long term. This conception was shaped by an imperfect analogy to bonds —the older, more familiar asset class against which equities were evaluated. Just as bond risk was understood primarily as the probability of default, equity risk was understood primarily as the probability that an investment would fail to return its principal in the long run. Graham and Dodd's *Security Analysis* (1934), a highly influential investment text of the era, exemplifies this view. Their framework distinguished between 'investments' —assets that promised 'safety of principal'— and 'speculations' —assets that did not. Risk, in this framework, was the possibility that a security would fail to be 'worth the price paid' under normal conditions: a conception that, while imprecise and difficult to operationalize, was fundamentally oriented toward adverse outcomes rather than the dispersion of potential outcomes.

In brief, during the three decades between Knight and Markowitz, academic economists referred to Knight's risk concept but did not operationalize it, and practitioners continued to work with their own loss-oriented conception. That is, they remained closer to the conventional understanding of risk as the possibility of adverse outcomes rather than to the dispersion-based conception that Knight had introduced.

It is also worth noting that the term 'volatility' —today the dominant synonym for risk in financial economics— was essentially absent from financial discourse during this period. Hartzmark and Solomon (2026) document that in the first decade of their newspaper sample, ending in 1915, the string ‘volatil*’ accounted for only 0.3% of all risk-related words appearing in the *Wall Street Journal and the New York Times* financial columns. As late as, 1951 —the year before Markowitz's paper— it still accounted for only 2% of risk-related words. The dominant risk-related terms during this period were ‘uncertainty,’ ‘fluctuat*,’ and ‘fear’ —qualitative, psychological, and directional in character, rather than statistical and symmetric. This linguistic evidence corroborates the conceptual picture: before Markowitz, neither the vocabulary nor the analytical framework of finance had converged on the standard deviation as a measure of risk.

In summary, the period between Knight (1921) and Markowitz (1952) was characterized by a curious paradox. Knight's redefinition of risk was widely acknowledged and frequently cited among academic economists, yet it produced no analytical tools and reshaped no empirical practice. Meanwhile, financial practitioners and financial writers continued to interpret risk in ways oriented toward adverse outcomes —closer in spirit to the pre-Knightian tradition than to Knight's own framework. It was only with Markowitz (1952) that Knight's definition was finally operationalized —paired with the standard deviation and embedded in a formal optimization framework. The nature of that pairing, and the independent problem it introduced, are examined in the following section.

The Knight-Markowitz Connection

In 1952, Harry Markowitz —who later was awarded the Nobel Prize in Economics— published his seminal paper *Portfolio Selection* (Markowitz, 1952), in which he addressed a practical problem of broad relevance to investors: given a fixed budget and a specific set of potential investment alternatives or assets, how should a rational investor determine the optimal allocation of that budget across available assets? In the financial economics literature, this is known as the optimal asset allocation problem.

Markowitz's framework rested on the assumption that it was possible to estimate the expected return of each asset. Specifically, he argued that historical data could be used to assign probabilities to different return scenarios and, consequently, calculate the expected return of each asset as the probability-weighted average of those scenarios. He further proposed that a rational investor should determine the optimal asset allocation—that is, the proportion of wealth invested in each asset—by solving an optimization problem that maximizes the portfolio's expected return subject to a constraint reflecting the investor's risk tolerance. But to formulate this optimization problem, he needed an explicit mathematical expression to measure risk.

It is worth noting that Markowitz's assumption regarding asset returns—a known set of potential outcomes, each associated with a known probability of occurrence—was entirely consistent with Knight's definition of risk: a situation of randomness in which the probabilities of all potential outcomes are known. Not surprisingly, Markowitz adopted the standard deviation of returns as the appropriate metric for risk assessment. Since the probability of each potential outcome was assumed to be known, the standard deviation could be easily calculated. In essence, he supplied what Knight's definition had conspicuously lacked: a precise formula for quantifying risk. The principal limitation of the standard deviation as a risk metric is that it fails to distinguish between favorable and unfavorable outcomes, treating positive and negative deviations from the mean as equivalent contributions to risk. This, of course, is not a deficiency of the standard deviation—which was designed to capture dispersion—but rather a natural consequence of Knight's risk definition. The following example clarifies this point.

Consider an investment that can result in three equally probable outcomes: (a) a return of 10%; (b) a return of 3%; and (c) a return of −1%. Suppose further that the investor requires a minimum return of 5%.

From this investor's perspective, risk is clearly associated with the realization of scenarios (b) or (c), i.e., outcomes falling below the 5% threshold. The expected return of this investment is 4%:

$$(1/3) \times (10\% + 3\% - 1\%) = 4\%.$$

The standard deviation, Markowitz's measure of risk, is 4.54%; it is simply the square root of:

$$(1/3) \times (10\% - 4\%)^2 + (1/3) \times (3\% - 4\%)^2 + (1/3) \times (-1\% - 4\%)^2.$$

Clearly, the metric adopted by Markowitz mixes both favorable and unfavorable scenarios from the investor’s viewpoint. Given that the investor requires a 5% return, a 10% outcome would be highly satisfactory, exceeding that threshold by a substantial margin. Nevertheless, within Markowitz's framework, this outcome is treated as a contributor to risk, solely on the grounds that it deviates from the mean, irrespective of direction. In brief, a risk measure that characterizes returns exceeding the investor's minimum requirement as 'undesirable' is at odds with both intuitive and practical conceptions of risk —and, more critically, it has the potential to lead to flawed investment decisions.

Why did Markowitz decide to use the standard deviation to measure risk? Did he have Knight’s definition of risk in mind? It is not possible to answer this question with any degree of certitude. Interestingly, Markowitz never made any explicit comment indicating that he was attempting to capture Knight’s concept with his decision. A plausible explanation is that he simply chose the path of least resistance through a difficult problem. It should be noted that the standard deviation already existed, having been introduced by Karl Pearson (Pearson, 1894, and was already widely used across the scientific community. In fact, during the first half of the twentieth century it had been employed extensively in quality control, psychometrics, biology, statistical inference, and related fields, although it was employed as a measure of dispersion rather than risk. More relevant in the context of Markowitz’s case, in economics, the variance (standard deviation squared) was deeply embedded in econometrics as a regular tool for measuring statistical uncertainty around parameter estimates. What had not been done was to extend the application of the standard deviation to assess risk in the context of portfolio selection problems, and more generally, in the context of financial decision-making. Given that Knight had introduced a conceptual definition of risk (admittedly, an unconventional one), but it came with no formula to make it operational, Markowitz’s important contribution was to connect these two existing ideas: he provided an established concept with the computational tool that it lacked. Whether this was intentional or not is largely irrelevant (Markowitz never mentioned Knight in his writings). What matters is that his decision gave operational force to an idea that until that moment was impossible to use in practice, that is, to carry out any calculations.

Moreover, Markowitz likely suspected that his broader ideas —the tradeoff between risk and return, the importance of diversification, and the introduction of the correlation matrix into investment decisions— would be controversial. And this is exactly what happened. The initial reaction of the financial and economic community to his ideas was lukewarm: the optimization-based framework he proposed relied on mathematical tools that were unfamiliar to most portfolio managers. Thus, the last thing he probably wanted was to introduce an entirely novel formula to estimate risk, since that would have only added to the perceived complexity of his framework. Combining Knight's notion of risk with the standard deviation —two pre-existing concepts— was simply a marriage of convenience, arguably, an easy decision unlikely to trigger too much resistance.

This assessment has been recently corroborated by Hartzmark and Solomon (2026). These authors, in *A Brief History of Financial Risk*, state that prior to Markowitz there is no record in the literature of the standard deviation of returns being used to assess risk. They attribute this, perhaps, to a lack of familiarity with the standard deviation among financial practitioners, to the computation being too burdensome, or simply to difficulties associated with obtaining the relevant data. Two authors who did use the standard deviation in their studies, Jackson (1928) and Cowles (1938), computed it for other purposes, not applied to market returns.

It is also worth noting that in 1952 —the same year as Markowitz's paper— A. D. Roy, a British economist, published *Safety First and the Holdings of Assets* (Roy, 1952), in which he independently proposed the standard deviation as a measure of risk, albeit without extending the analysis to a comprehensive treatment of the portfolio selection problem. Strictly speaking, Roy also deserves at least partial recognition for introducing the standard deviation as a risk metric. Whether that credit belongs to Markowitz alone or should be shared with Roy is, however, immaterial to our discussion. The key point is that the standard deviation emerged as a risk metric in 1952.

Finally, it is important to separate Markowitz's metric choice from his broader intellectual contributions, which were substantial and largely independent of it. He was the first to cast the portfolio selection problem within a rigorous quantitative framework. His central insights—the importance of diversification and the necessity of accepting higher risk in pursuit of higher returns—have withstood the test of time. They gave rise to the field now known as modern portfolio theory and established a principle important in its own right: that risk is a quantity that can, and should, be measured formally. The limitations identified in this paper concern specifically the instrument Markowitz chose to measure risk—not the decision to measure it, nor the broader theoretical structure he built around it.

Finance post-Markowitz

Notwithstanding the practical challenges associated with the implementation of Markowitz's framework (e.g., its tendency to produce what in mathematical optimization is known as corner solutions, or, unstable asset allocation percentages) one of its key contributions was rapidly embraced by economics: the use of the standard deviation as a proxy for risk —a proxy that carried embedded within it both Knight's conceptual framework and the limitations that framework imposed.

Shortly after Markowitz's seminal work, Tobin (1958) incorporated the standard deviation of returns into his definition of utility to capture investors' risk aversion. Sharpe (1964) subsequently used the standard deviation of returns to develop the Capital Asset Pricing Model (CAPM), and two years later, to introduce the Sharpe ratio, a measure of risk-adjusted returns (Sharpe, 1966). However, the definitive consolidation of the standard deviation of returns as a formal risk metric is often attributed to Black and Scholes (1973). In deriving their option valuation formulas, they adopted the term volatility to refer to the standard deviation of returns. This terminology proved widely popular, and today, volatility is the preferred term among finance and economics professionals to refer to the standard deviation.

Nevertheless, it is important to recognize that a sober assessment of these models is not encouraging. To be more precise, although these models have been an academic success, their practical application has been problematic. They are part of the curriculum in most economics and finance courses —probably because they are ideally suited to structuring easy-to-grade student assignments, and because the mathematics behind them is both impressive and aesthetically attractive. Additionally, these models have generated a copious amount of academic literature. Yet they have failed where it matters most: financial practitioners do not use them, or at best employ heavily modified versions that bear little resemblance to the original in terms of complexity or structure. This lack of practical enthusiasm can be traced to a single root cause: most of their results have not been supported by empirical evidence. This is precisely what Mäki (2018) warned against in *The Rights and Wrongs of Economic Modelling*: the danger of relying on assumptions that, while making models mathematically tractable, are too simplistic to capture the nuances of reality.

A case in point is Fama and French's evaluation of the CAPM's validity —arguably the most exhaustive assessment of this model to date, and a damning one at that. Their assessment is worth quoting: ‘The attraction of the CAPM is that it offers powerful and intuitively pleasing predictions about how to measure risk and the relation between expected return and risk. Unfortunately, the empirical record of the model is poor,’ (Fama and French, 2004).

The verdict regarding the Black and Scholes valuation formulas is not better. The empirical evidence shows that the prices predicted by this model are at odds with market reality (e.g., Redroban and Cifuentes, 2021). Additionally, there is plenty of evidence that traders do not use the Black and Scholes model to price the options they buy and sell (Haug and Taleb, 2011).

These examples are very telling because they refer to situations in which getting ‘wrong’ results —since often substantial amounts of money are involved— can lead to financial ruin or heavy losses. In short, unlike some academic debates about arcane topics which are very seldom settled, financial mistakes can have immediate and very palpable consequences. Therefore, the fact that people who have ‘skin in the game' (a colloquial term used by those whose financial decisions have direct personal consequences) do not use these models is noteworthy.

It might be argued that the above-mentioned criticism is unfairly harsh, and that practitioners do in fact use some of these models. The Black-and-Scholes model for pricing derivatives, for instance, is frequently employed to construct hedging strategies —a purpose quite distinct from its original intent— or to extract implied volatilities, an application that is in fact inconsistent with one of the model's own basic assumptions. Nevertheless, the relevant fact is that the Black-and-Scholes model does not work for the purpose it was conceived and it is not used for the purpose it was intended. This topic has been discussed in detail elsewhere (Weatherall, 2018; Cifuentes and Charlin, 2024).

In fairness, it should also be recognized that the models mentioned above, in addition to relying on the standard deviation to assess risk, were also built upon a host of other assumptions that have proved problematic. For example, the normality of asset returns, symmetric risk-aversion preferences, and the adequacy of Brownian motion to mimic the time evolution of financial variables, among others. No doubt, these assumptions have also contributed to the poor performance of such models, and therefore it would be a stretch to place all the responsibility on the standard deviation alone. That said, the fact that the standard deviation is a structural component common to all of these models should not be lightly dismissed.

It is worth noting, however, at the risk of sounding repetitive, that the problem is not related to an intrinsic shortcoming of the standard deviation. Rather, it is due to the fact that the standard deviation happens to capture the key features of a flawed (Knightian) definition of risk. In other words, we are simply applying the right tool to interpret a faulty idea.

The Economics-Finance Divergence

Toward the end of the twentieth century, there was an important and consequential innovation: Tyrrell Rockafellar and Stanislav Uryasev introduced a risk metric they called Conditional-Value-at-Risk or CVaR (Rockafellar and Uryasev, 2000). The CVaR —also known as Expected Shortfall— measures risk by focusing exclusively on adverse outcomes. That is, it is fully compatible with the conventional understanding of risk as the possibility of harm or loss. In this respect, it represented a departure from the standard deviation. In short, it not only embraced a notion of risk compatible with common usage, but also signaled a return to the notion of risk that prevailed among economists before Knight. The CVaR represented a correction at two levels simultaneously: not only did it embody a conception of risk based on the possibility of experiencing adverse outcomes, but it also provided a metric that operationalized that conception faithfully.

The CVaR also offers two attractive properties. First, it satisfies an important condition known as subadditivity, that is, the CVaR is coherent. A full discussion of this mathematical property is beyond the scope of this paper, but it suffices to say that the CVaR does not penalize diversification. This is important since diversification is not only a desirable feature highly sought by investors, but also, it is one of the key concepts whose relevance Markowitz emphasized. Second, the CVaR captures not only investors' attitude toward risk, but people's attitude toward risk more broadly. In fact, several years before the CVaR was formulated, Kahneman and Tversky (1979) established through experimental research that people's responses to losses are asymmetrical: they fear losses far more than they value gains. This asymmetric attitude toward risk is easily captured by the CVaR (it totally ignores gains); whereas the standard deviation, by construction, treats gains and losses symmetrically. A technical aside: the fact that return distributions are not normal —that is, they are highly asymmetrical (pronounced skewness) and exhibit heavy tails— only exacerbates the shortcomings of the standard deviation as a risk measure.

The introduction of the CVaR had two significant consequences. On the practical and operational side, it was rapidly adopted by financial practitioners, rather than by the academic community, including portfolio managers and risk managers in banks, insurance companies, and pension funds. And its acceptance was also extended to major regulatory frameworks; for example, Basel III (2010), which governs risk management in banking institutions, relies explicitly on the CVaR, as does the Swiss Solvency Test (2006), the regulatory standard applied to insurance companies. Moreover, at the conceptual level, the emergence of the CVaR marked a clear and radical separation between the type of financial economics taught in universities and the type of tools employed in the then nascent, but now more mature, field known as financial engineering.

The conceptual divide between conventional economics and financial engineering in their respective conceptions of risk has now become firmly entrenched. Economics continues to embrace a notion of risk shaped largely by Knight's work, relying on the standard deviation as the *de facto* metric for its quantification. The enduring appeal of the standard deviation is, in no small part, attributable to its mathematical tractability: there is a rich body of analytical formulas as well as many theoretical results based on this metric.

Financial engineering, by contrast, has moved increasingly toward equating risk with losses and adverse outcomes, adopting the Conditional-Value-at-Risk (CVaR) as its preferred risk assessment metric, often complemented by additional loss-based measures such as maximum drawdown and tail-value-at-risk. It might be objected that the CVaR lacks, at least computationally, the elegance of the standard deviation, and that its definition —at least its continuous version— is comparatively less intuitive and more mathematically challenging.

The discrete formulation of the CVaR, however, which is in practice the version most widely employed in applied settings, is straightforward to implement and easily communicable to non-specialist audiences. Moreover, a significant advantage of the CVaR in the context of portfolio optimization is that it satisfies a convexity condition, which, in turn, lends itself to a linear programming formulation that can be solved efficiently using most standard commercial optimization software. This computational convenience has contributed substantially to its growing adoption among practitioners and has reinforced the methodological divergence between these two disciplines.

Conclusion

Economics is a social science. It is largely concerned with studying the way individuals, firms, and government agencies choose among several alternatives when resources are limited and trade-offs are unavoidable. Financial economics, in turn, is more concerned with investments, and more precisely, with making investment decisions under conditions of uncertainty. Evidently, many such decisions entail the possibility of adverse outcomes, and thus, a rigorous conceptualization of risk —including both its definition and measurement— becomes essential to any framework aimed at developing analytical models to facilitate, or to understand, such processes. Irrespective of the debates about rationality or utility maximization, an adequate risk characterization —along with a suitable risk metric— constitutes a fundamental component of any decision-making model. Moreover, a flawed characterization of risk will inevitably produce models with poor explanatory or predictive power; in other words, models of limited practical use. Worse yet, such models can lead to outright poor decisions.

This last point is critical. The financial (subprime mortgage) crisis of 2007/2008 has been almost universally described as a financial risk management failure (see for example, Senior Supervisors Group, 2009; Sacasa, 2008; Ashby, 2010). It would certainly be an overstatement to assert that faulty risk models —models based on the standard deviation as a risk metric— triggered this financial crisis. But the fact that most risk models used at that time were based on volatility is telling. While these models did not cause the crisis, they surely exacerbated it by providing a misleading risk assessment. In short, notwithstanding the relevance of having conceptual clarity (a theoretical issue) in terms of what risk means, it must be recognized that relying on a faulty characterization of risk, that is, focusing on dispersion of potential outcomes (volatility) and not losses, can also have dire practical consequences.

Prior to Knight's influential book, economics had a coherent understanding of risk, as evidenced in the works of Smith, Ricardo, and other classical economists. Their conception of risk was in line with interpretations prevalent in other disciplines and in common usage. Hence, Knight's redefinition of risk, rather than addressing a conceptual confusion, set economics on a problematic path—one that was later given operational form by Markowitz's adoption of the standard deviation as a risk metric. But this metric (notwithstanding its adequacy for many other purposes) is inconsistent with any meaningful concept of risk as the possibility of adverse outcomes, since it focuses on the dispersion of the distribution rather than solely on the side of the distribution describing losses. Knight's flawed definition, made operationally powerful by Markowitz's adoption of a metric compatible with it, resulted in an unfortunate yet powerful combination—one that gave rise to a succession of models whose empirical validation has been consistently weak at best.

Nevertheless, it would be unfair not to acknowledge Knight's book's main contribution: a compelling argument to explain why companies earn profits in a free market economy. His elaborations on risk and return were a secondary focus, rather than the central topic of his work. The shortcomings of the standard deviation as a risk metric—its failure to distinguish between favorable and adverse outcomes—should not be taken as an invitation to abandon it as an analytical tool, but rather as a signal that it has been applied to the wrong concept. Indeed, the standard deviation, as a measure of dispersion, is precisely the right instrument for measuring uncertainty—understood in its conventional sense as incomplete knowledge about future outcomes, irrespective of their direction. In fact, when the standard deviation was first introduced, that was precisely the intent: to capture dispersion, that is, uncertainty.

The analysis presented here points to an inevitable recommendation: that economics would benefit from realigning its conceptualization of risk with the approach long adopted in financial engineering —specifically, by returning to the pre-Knightian view of risk as the possibility of experiencing undesirable outcomes. Reframing risk in this manner would allow the discipline to correct a long-standing conceptual misalignment. After more than a century of empirical evidence and theoretical developments (Knight’s work dates to 1921), a reassessment of Knight’s distinction between risk and uncertainty is both appropriate and overdue. The verdict is hard to escape: Knight's definition of risk has generated neither significant theoretical insights nor practical benefits for economics.

References

Albarrán, P., Alonso, J., Herrero, C., Ponti, G., Sartarelli, M., & Topazio, D. (2025). An experiment on outcome uncertainty. *Journal of Risk and Uncertainty*, *70*, 147–170. https://doi.org/10.1007/s11166-024-09448-0

Ashby, S. (2010). The 2007–2009 financial crisis: Learning the risk management lessons. Nottingham: Financial Services Research Forum. https://www.nottingham.ac.uk/business/who-we-are/centres-and-institutes/gcbfi/documents/researchreports/paper65.pdf

Ayyub, B. M. (2015). Introduction to the aims and scope of the journal. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part B: Mechanical Engineering*, *1*(1), 1–4. https://doi.org/10.1061/AJRUA6.0000001

Bahr, R., & Holme, I. (2003). Risk factors for sports injuries—a methodological approach. *British Journal of Sports Medicine*, *37*(5), 384–392. https://doi.org/10.1136/bjsm.37.5.384

Barratt, B. I. P., Meenken, E. D., & Withers, T. M. (2023). Characterising uncertainty in risk assessments for biological control: Using case studies from New Zealand to inform future research. *BioControl*, *68*, 101–115. https://doi.org/10.1007/s10526-023-10181-x

Basel Committee on Banking Supervision. (2010). Basel III: A global regulatory framework for more resilient banks and banking systems. Bank for International Settlements. https://www.bis.org/publ/bcbs189.htm

Bewley, T. F. (2002). Knightian decision theory: Part I. *Decisions in Economics and Finance*, *25*(2), 79–110. (Original work circulated as Cowles Foundation Discussion Paper No. 807, 1986) https://doi.org/10.1007/s102030200006

Bhide, A. (2021). Renewing Knightian uncertainty: Problems and possibilities. Harvard Business School Research Paper Series No. 21-129. https://ssrn.com/abstract=3868219

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, *81*(3), 637–654. https://doi.org/10.1086/260062

Botticini, M., Buri, P., & Marinacci, M. (2023). Presidential address 2023: The beauty of uncertainty: The rise of insurance contracts and markets in medieval Europe. *Journal of the European Economic Association*, *21*(6), 2287–2326. https://doi.org/10.1093/jeea/jvad059

Cantillon, R. (1755). Essai sur la nature du commerce en général [Essay on the nature of trade in general] (H. Higgs, Trans.). Macmillan. (Original work published 1755)

Cifuentes, A., & Charlin, V. (2024). Further reflections based on the Black–Scholes (B-S) model. *Philosophy of Science*, 1–6. https://doi.org/10.1017/psa.2024.31

Cowan, D. (2016). Knightian uncertainty. In *Frank H. Knight* (Great Thinkers in Economics, pp. 27–74). Palgrave Macmillan. https://doi.org/10.1057/978-1-137-46211-4_2

Cowles, A., 3rd. (1938). Common stock indexes, 1871–1937. Principia Press. (Cowles Commission Monograph No. 3)

Cox, L. A. (2012). Confronting deep uncertainties in risk analysis. *Risk Analysis*, *32*(10), 1607–1629. https://doi.org/10.1111/j.1539-6924.2012.01792.x

Duke, L. D., & Taggart, M. (2000). Uncertainty factors in screening ecological risk assessments. *Environmental Toxicology and Chemistry*, *19*(6), 1668–1680. https://doi.org/10.1002/etc.5620190626

Ellingwood, B. R., & Kinali, K. (2009). Quantifying and communicating uncertainty in seismic risk assessment. *Structural Safety*, *31*(2), 179–187. https://doi.org/10.1016/j.strusafe.2008.06.001

Faber, M. H., & Stewart, M. G. (2003). Risk assessment for civil engineering facilities: Critical overview and discussion. *Reliability Engineering & System Safety*, *80*(2), 173–184. https://doi.org/10.1016/S0951-8320(03)00027-9

Fama, E. F., & French, K. R. (2004). The capital asset pricing model: Theory and evidence. *Journal of Economic Perspectives*, *18*(3), 25–46. https://doi.org/10.1257/0895330042162430

Graham, B., & Dodd, D. L. (1934). Security analysis. Whittlesey House, McGraw-Hill.

Hansson, S. O. (2004). Philosophical perspectives on risk. *Techté: Research in Philosophy and Technology*, *8*(1), 10–35. https://doi.org/10.5840/techne2004818

Hansson, S. O. (2023). Risk. In E. N. Zalta & U. Nodelman (Eds.), *Stanford encyclopedia of philosophy* (Summer 2023 ed.). https://plato.stanford.edu/archives/sum2023/entries/risk/

Hartzmark, S. M., & Solomon, D. H. (2026). A brief history of financial risk. [Working paper]. https://ssrn.com/abstract=6991358

Haug, E. G., & Taleb, N. N. (2011). Option traders use (very) sophisticated heuristics, never the Black–Scholes–Merton formula. *Journal of Economic Behavior & Organization*, *77*(2), 97–106. https://doi.org/10.1016/j.jebo.2010.09.013

Hausman, D. M. (2023). The inexact and separate science of economics (2nd ed.). Cambridge University Press.

Jackson, J. R. (1928). Common and preferred stocks as investments. *Journal of Business of the University of Chicago*, *1*(3), 294–323. https://www.jstor.org/stable/2349530

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291. https://doi.org/10.2307/1914185

Knight, F. H. (1921). Risk, uncertainty and profit. Houghton Mifflin.

LeRoy, S. F., & Singell, L. D., Jr. (1987). Knight on risk and uncertainty. *Journal of Political Economy*, *95*(2), 394–406. https://doi.org/10.1086/261461

Makower, H., & Marschak, J. (1938). Assets, prices and monetary theory. *Economica*, *5*(19), 261–288. Reprinted in G. J. Stigler & K. Boulding (Eds.), *Readings in price theory* (pp. 283–328). Richard D. Irwin for the American Economic Association, 1952.

Mäki, U. (2018). Rights and wrongs of economic modelling: Refining Rodrik. *Journal of Economic Methodology*, *25*(3), 216–226. https://doi.org/10.1080/1350178X.2018.1488475

Markowitz, H. M. (1952). Portfolio selection. *Journal of Finance*, *7*(1), 77–91. https://doi.org/10.1111/j.1540-6261.1952.tb01525.x

Marschak, J. (1938). Money and the theory of assets. *Econometrica*, *6*(4), 311–325. Reprinted in D. Cass & L. W. McKenzie (Eds.), *Selected readings in macroeconomics and capital theory from Econometrica* (pp. 74–88). MIT Press, 1974.

Mill, J. S. (1870). Principles of political economy (7th ed.). Longmans, Green, Reader and Dyer.

Mitchell, W. C. (1922). Review of Risk, uncertainty and profit by F. H. Knight. *American Economic Review*, *12*(2), 274–275. https://www.jstor.org/stable/1802628

Packard, M. D., Bylund, P. L., & Clark, B. B. (2021). Keynes and Knight on uncertainty: Peas in a pod or chalk and cheese? *Cambridge Journal of Economics*, *45*(5), 1099–1125. https://doi.org/10.1093/cje/beab025

Pearson, K. (1894). Contributions to the mathematical theory of evolution. *Philosophical Transactions of the Royal Society of London A*, *185*, 71–110.

Redroban, S., & Cifuentes, A. (2021). On the performance of the Black and Scholes (B-S) options pricing formulas during the subprime and Covid-19 crises. *Journal of Corporate Accounting & Finance*, *32*(3), 75–85. https://doi.org/10.1002/jcaf.22504

Ricardo, D. (1817). On the principles of political economy and taxation. John Murray.

Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, *2*(3), 21–41. (Working paper version circulated 1999) https://doi.org/10.21314/JOR.2000.038

Roy, A. D. (1952). Safety first and the holdings of assets. *Econometrica*, *20*(3), 431–449. https://doi.org/10.2307/1907413

Sacasa, N. (2008). Preventing future crises: Priorities for regulatory reform after the meltdown. *Finance and Development*, *45*(4), 11–14. International Monetary Fund. https://www.imf.org/external/pubs/ft/fandd/2008/12/pdf/sacasa.pdf

Scherer, F. M. (2007). Uncertainty and choice: The challenges of pharmaceutical efficacy, safety, and cost. *Managerial and Decision Economics*, *28*(4–5), 373–385. https://doi.org/10.1002/mde.1340

Schmalensee, R. (1974). Consumer behavior versus economic theory. *Recherches Économiques de Louvain / Louvain Economic Review*, *40*(3), 261–276. https://www.jstor.org/stable/40723430

Senior Supervisors Group. (2009). Risk management lessons from the global banking crisis of 2008. https://www.sec.gov/news/press/2009/report102109.pdf

Shackle, G. L. S. (1949). Expectation in economics. Cambridge University Press.

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, *19*(3), 425–442. https://doi.org/10.1111/j.1540-6261.1964.tb02865.x

Sharpe, W. F. (1966). Mutual fund performance. *Journal of Business*, *39*(1), 119–138. https://doi.org/10.1086/294846

Smith, A. (1776). An inquiry into the nature and causes of the wealth of nations. W. Strahan and T. Cadell.

Swiss Financial Market Supervisory Authority (FINMA). (2006). Swiss solvency test. FINMA.

Tobin, J. (1958). Liquidity preference as behavior towards risk. *Review of Economic Studies*, *25*(2), 65–86. https://doi.org/10.2307/2296205

Turvey, R., de Graaff, J., Baumol, W. J., & Shackle, G. L. S. (1949). Three notes on Expectation in Economics. *Economica*, *16*(64), 336–346. https://doi.org/10.2307/2549588

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131. https://doi.org/10.1126/science.185.4157.1124

Watkins, G. P. (1922). Knight's 'Risk, uncertainty and profit.' *Quarterly Journal of Economics*, *36*(4), 682–690. https://doi.org/10.2307/1884757

Weatherall, J. O. (2018). The peculiar logic of the Black-Scholes model. *Philosophy of Science*, *85*(5), 1152–1163. https://doi.org/10.1086/699677

Wilks, D. S., & Horowitz, K. A. (2014). A novel financial market for mitigating hurricane risk. Part I: Market structure and model results. *Weather, Climate, and Society*, *6*(3), 257–270. https://doi.org/10.1175/WCAS-D-13-00032.1

Wood, O. G. (1964). Evolution of the concept of risk. *Journal of Risk and Insurance*, *31*(1), 83–91. https://doi.org/10.2307/251211

APPENDIX: The Evolution of the Concept of Risk

From common usage, through Knight and Markowitz, to its correction in financial engineering

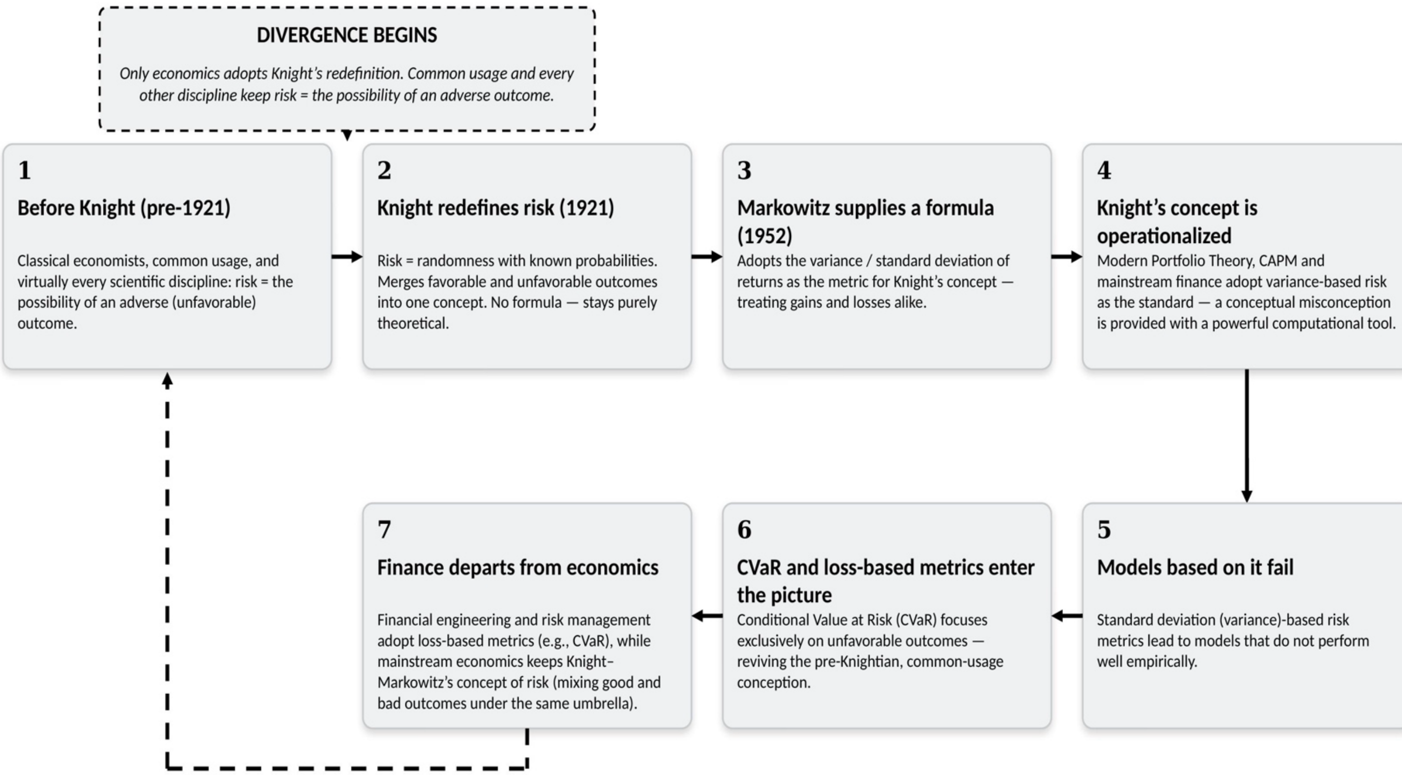

Finance and financial engineering return to the pre-Knightian / common-usage conception of risk (loss-based).

Centro UC
CLAPES UC
Centro Latinoamericano de Políticas Económicas y Sociales

clapesuc

@clapesuc

@claapes_uc

Claapes UC

ClapesUC

www.clapesuc.cl