

# Mariam Thalos

---

## Expectational v. Instrumental Reasoning: What Statistics Contributes to Practical Reasoning

---

Diametros nr 53, 125-149

---

2017

Artykuł został opracowany do udostępnienia w internecie przez Muzeum Historii Polski w ramach prac podejmowanych na rzecz zapewnienia otwartego, powszechnego i trwałego dostępu do polskiego dorobku naukowego i kulturalnego. Artykuł jest umieszczony w kolekcji cyfrowej [bazhum.muzhp.pl](http://bazhum.muzhp.pl), gromadzącej zawartość polskich czasopism humanistycznych i społecznych.

Tekst jest udostępniony do wykorzystania w ramach dozwolonego użytku.

# EXPECTATIONAL V. INSTRUMENTAL REASONING: WHAT STATISTICS CONTRIBUTES TO PRACTICAL REASONING

– Mariam Thalos –

**Abstract.** Utility theories—both Expected Utility (EU) and non-Expected Utility (non-EU) theories—offer numericalized representations of classical principles meant for the regulation of choice under conditions of risk—a type of formal representation that reduces the representation of risk to a single number. I shall refer to these as *risk-numericalizing theories of decision*. I shall argue that risk-numericalizing theories (referring both to the representations and to the underlying axioms that render numericalization possible) are not satisfactory answers to the question: “How do I take the (best) means to my ends?” In other words, they are inadequate or incomplete as instrumental theories. They are inadequate because they are poor answers to the question of what it is for an option to be *instrumental* towards an end. To say it another way, they do not offer a sufficiently rich account of what it is for something to be a means (an instrument) toward an end.

**Keywords:** statistics; distributions; risk; instrumental reasoning; expectational reasoning.

## 1. Introduction

Classical decision theory is considered, at least by some, the best-worked-out articulation of the idea that practical reason is thoroughly instrumental. This basic philosophical idea is expressed perhaps most forcefully in Bertrand Russell’s dictum that “rational choice” signifies choice of the right means to your ends, and has nothing whatever to do with the choice of ends themselves. The idea is regarded as so fundamental as to require no defense. (So, for instance, Lara Buchak simply takes it for granted that decision theory is the formalization of means-ends reasoning, and on the strength of the assumption raises objections to certain defenses of canonical Expected Utility theory mounted against some classic challenges.<sup>1</sup>) Accordingly, we can refer to this idea as *instrumentalism*.

Within classical decision theory, instrumental reasoning is also an *all-things-considered* enterprise, concerned not merely with the realization of individual ends, but also with coordinating achievement of their sum in the relevant context. The

---

<sup>1</sup> Buchak (2014).

fundamentals of how to advance a rich menu of potentially conflicting concerns simultaneously, in the relevant context, is the sole focus of decision theory. The most celebrated solution to this problem is the cornerstone of contemporary decision science: the very notion of *utility*. The concept of a totality of ranked concerns is at the core of the concept of utility. The theory built upon the concept – the theory of Expected Utility (EU) – is meant to give highest expression to the maximizing conception of instrumental rationality, according to which the agent’s fundamental task is to choose, from among all “moves” available, the one that best advances the totality of his or her concerns into the indefinite future.

Instrumentalism has been challenged. Some prominent challengers have contended that there is – as there must be – room as well for reasoning about the *adoption* of goals or concerns, within the core of practical reasoning. Practical reasoning, say these challengers, must also regulate the adoption of goals, not merely work out how best to advance the sum of whichever happen to have been adopted. Other challengers suggest that articulation of options, in standard applications of classical decision theory, already introduces an element of value, so is not entirely value-neutral as its apologists purport. Moreover, in practical life, “frames” render the service of articulating options, and since EU as routinely put into practice employs no “checks” to ensure that options are framed so as to assure alignment with the agent’s values, EU is (once again) not value-neutral.<sup>2</sup> Still other challengers purport that classical decision theory is ham-fisted, handling inadequately the demands of instrumental reasoning in exceptional risk conditions – for instance where probabilities are very small or are in some other way unusual. Hampton can be read this way,<sup>3</sup> as can many of the architects of the different varieties of non-expected utility (non-EU) theories.<sup>4</sup> The view I shall be advancing here is a variation on this last position: EU, very much like its successors in the decision-theoretic literature, is too ham-fisted to be legitimately viewed as theory of instrumental rationality.

Utility theories (both EU and non-EU) offer numericalized representations of classical principles meant to regulate choice – a specific type of formal representation that reduces the representation of risk to a single number. I shall refer to these as *risk-numericalizing theories of decision*. I shall argue that risk-numericalizing

---

<sup>2</sup> A “frame” is a set of beliefs, attitudes, values, mental models, etc., that is utilized in the apprehension of a decision situation, as a lens of interpretation. The literature on the framing of choice situations is now vast, but the locus classicus is Tversky and Kahneman (1981). See also Tversky, Kahneman (1986).

<sup>3</sup> Hampton (1994).

<sup>4</sup> E.g. Machina (1982); Quiggin (1982), as well as philosophical work such as Buchak (2014).

theories (referring to the representations or to the underlying axioms that render numericalization possible) are not satisfactory answers to the question: “How do I take the (best) means to my ends?” In other words, they are inadequate and/or incomplete as instrumental theories, and so are simply not full-throated theories of instrumental reasoning. In other words, that they are not adequately answering the question: how do I take the (best) means to my ends? They miss out on too much that bears on that question.<sup>5</sup> So they are ham-fisted as theories of instrumental reasoning. I shall argue that the demands of instrumental reasoning are subtler and more demanding than is usually thought.

The reason for the divergence between instrumental reasoning and risk-numericalizing reasoning is fundamentally this: risk-numericalizing reasoning is concerned with weighing straight utility in such a way that an option has the numerical value it does in relation to only two (albeit important) categories of considerations: (1) our present goals all considered; and (2) all potential opportunities for present action. One might think (and not unreasonably) that this makes risk-numericalizing reasoning a very admirable rendering of instrumental reasoning. But an option, in the risk-numericalizing framework, is not evaluated—not numerically nor in any other way—in relation to *other potential decisions*. Decisions, in other words, are not automatically construed as made in the context of other decisions of the same agent—instead, they’re treated as *independent*. And this means that *future* potential opportunities are not taken into consideration. In other words, numericalized theories of the sort that have become familiar evaluate only the (single) *present* action, but do not take account of how the present action *aggregates with other actions that follow on from earlier decisions, forming arcs through an agent’s global option space over a lifetime or a substantial portion thereof*. In other words, risk-numericalizing analyses fail to consider the way that decisions, as such, scale over an agent’s lifetime. From a global perspective, however, decisions are positively not independent; and from a normative perspective they indeed ought not to be. Because in fact decisions are themselves *networks*, over a rich stream of *opportunities* whose constitution is itself impacted by agents’ activities. Consequently, the agent who reaps rewards from actions as the rewards come in (and not merely at the time of action) must enjoy or suffer them as an aggregate.

This reality, that risk-numericalizing theories remain mired in a construal of decisions as taken individually, is symptomatic of a greater failing—a failing to

---

<sup>5</sup> The contrast between instrumental reasoning and the reasoning implicit in risk-numericalizing theories is perhaps starker still in contexts of nonnormal statistical distributions. Cf. Thalos, Richardson (2013).

appreciate that choices have to be justified as instruments *collectively* serving the agent's *collective* goals and concerns. The primary aim of this essay is to demonstrate that EU falls short in this regard, and so falls short as a theory of instrumental reasoning. A theory of instrumental reasoning should present us with a complete account in which the measure of a decision is taken in terms of how the decision comports with other decision in the agent's portfolio of reasons and actions, to serve as instruments for the advancement of the aims of an agent whose goals and concerns are themselves networked and dynamic over time. For instance, the value of a decision to brush one's teeth on a given occasion is networked with plans, habits, and protocols already in place regarding the brushing of teeth, so that someone who plans on skipping a night of brushing would have to consider the impact of that on future brushings. And so it is no wonder, as we will see, that risk-numericalizing reasoning cannot capture the full richness of instrumental reasoning, and in quite ordinary circumstances.

But this essay will make good on these contentions through examples that involve only money. We will thus avoid controversy, and avoid too the appearance of cheating or special pleading.

## **2. The principle of bundling**

In a famous article Paul Samuelson reports that he once offered a colleague the favorable side of a win \$200/lose \$100 wager on a fair coin toss.<sup>6</sup> The colleague (now commonly referred to as "Samuelson's Colleague" - SC) declined the bet, but declared a willingness to accept 100 such bets together. Samuelson purported to show that this pair of choices was inconsistent with expected utility theory, to the effect that if, for a relatively constant range of wealth, a person turns down a particular gamble, then the person should also turn down an offer of a string of such gambles. "Each outcome must have its utility reckoned at the appropriate probability; and when this is done it will be found that no sequence is acceptable if each of its single plays is not acceptable. This is a basic theorem." To keep on the right side of EU theory, he maintained, one must have approximately the same risk attitude towards an aggregation of independent, identical gambles, as towards each of the gambles taken singly.

Samuelson was very careful to argue that this principle is independent of risk aversion—hence the restriction to a fixed wealth level. And he was careful too to emphasize that this principle is not a demand for risk neutrality; it is instead, he insisted, a demand that one exercise consistency in one's aversion to risk. But what

---

<sup>6</sup> Samuelson (1963).

kind of consistency is this? That is our question. And is this sort of consistency worth having? We will answer No. Samuelson's consistency amounts to a certain constraint of independence upon an agent's decisions. And indeed we should embrace the constraint if we are obliged to treat decisions taken by the same individual over a given interval as necessarily *unlinked*. But why should we do that? It is not wisdom.

### 3. Great expectations: Some preliminaries

Suppose an agent is deciding among a slate of potential actions  $\{A_1, A_2, \dots, A_p\}$ . Let's refer to a possible action as a *prospect*. Suppose that if prospect  $A_j$  is selected, the outcomes that can result are some  $n$  in number  $(O_{j1}, \dots, O_{jn})$ , which the agent assesses differentially at utilities  $u(O_{j1}), \dots, u(O_{jn})$ . Assume the probabilities associated with each outcome are specified:  $P(O_{ji}) = p_{ji}$ , such that  $\sum_{i=1}^n p_{ji} = 1$ . If the probabilities are not trivial (not all either zeros or ones), then our agent is facing a *decision under risk*. Let's designate by the term *expectation of  $A_j$*  –  $E(A_j)$  for short – the quantity computed as follows:

$$E(A_j) = \sum_{i=1}^n p_{ji} \times u(O_{ji}).$$

The fundamental theorem of Expected Utility (EU) dictates that our agent should assess the utility of prospect  $A_j$  by its expectation  $E(A_j)$ , and make their choice accordingly: perform that action  $A_j$ , of those actions available now  $\{A_1, A_2, \dots, A_p\}$ , with the highest expected utility. Let's refer to the directive to decide on the basis of expectations as the *expectation principle*. This principle is meant to capture, in risk-numericalized terms, the idea that picking the  $A_j$  with the largest EU is the agent's way of picking the best means to her ends.<sup>7</sup> It will be my contention that the two concepts are quite distinct – the  $A_j$  with the largest EU and the best means to one's ends frequently diverge.

To begin, let's consider how EU defends the expectation principle.<sup>8</sup> The expectation principle is a *theorem* of the theory; it follows from assumptions that are very abstract, and purportedly more elemental than the theorem itself. Put together, these handful of elemental assumptions issue in the expectation principle. There exist a number of formulations of EU (and hence of the expectation princi-

---

<sup>7</sup> In effect, it directs agents to constrain the way that they generate their attitudes toward prospects, based on their (uncriticizable, but nonetheless mandated) attitudes towards their component outcomes, via a certain rule of combination – namely, via an expectational rule.

<sup>8</sup> There are arguments external to EU itself in favor of this constraint as well – arguments that I cannot address here: a de Finetti inspired argument appealing to the rationality of working with best estimates [(Pettigrew (2015))] and a school of thinking on the topic of sequential reasoning discussed most prominently by McClennen (1990).

ple). But built into each such way of formulating elemental assumptions is the fundamental idea that every choice is a gamble, that gambles are evaluated using principles of probability theory, and that these gambles can be ordered in a continuous way. So, for example, the von Neumann-Morgenstern formulation employs a *continuity condition* that has the following immediate consequence:

If P, Q and R are prizes such that P is preferred to Q, which in its turn is preferred to R, then there exists a real number  $0 \leq \alpha \leq 1$ , such that a gamble with prizes P (with probability  $p = \alpha$ ) and R (with probability  $q = 1 - \alpha$ ) that is assessed as equivalent to the prize Q.

It's easy to see that this condition has application in the case of Samuelson's colleague SC. Let  $P = \$200(\text{gain}) \times 100$ ,  $Q = 0$ , and  $R = \$100(\text{loss}) \times 100$ . Then there will be odds at which SC will be willing to accept a wager (or a series of such wagers) that involves P and R as "prizes" (and intermediate prizes as well, if we are talking about the series). According to EU's assumptions, there *must* be such odds! The assumptions of EU are such that there can be nothing – nothing at all – we can withhold as stakes.

The question can now be raised whether the colleague should be required to accept each of the wagers (single and aggregate) at the *same* odds. To be sure, EU does not require SC to find acceptable tomorrow what he finds acceptable today – so long as he changes other preferences accordingly. (In other words, the basic principles of EU are not sensitive to the timing of attitudes toward benefits and burdens, and all decisions are treated as static therein; extensions of the theory to take account of consumption as time-indexed have been controversial.<sup>9</sup>) When SC is compliant with the precepts of EU at each moment of decision, we can say he is *synchronically* in compliance with the precepts of EU. The point to note, therefore, is that EU makes no *diachronic* demands on its patrons; as they travel through time, agents can completely transform their preferences, so long as at any instant those preferences conform to the synchronic axioms of coherence. Similarly, EU does not insist – and neither does Samuelson or his fellow travelers Tversky and Bar-Hillel – that the colleague's willingness to accept a wager be independent of his level of wealth.<sup>10</sup> His decision to accept can vary depending on how much money he currently has to his name. And so EU allows its patrons to be averse to taking big risks with current holdings – SC is allowed, in other words, to demand very favorable odds before staking something currently in his possession.

---

<sup>9</sup> Frederick, Loewenstein, Donoghue (2002).

<sup>10</sup> Tversky, Bar-Hillel (1983).

In light of all this freedom for risk aversion (laxity in the theory, if you will), how is SC going amiss, according to Samuelson? Let us first show that Samuelson's "proof" of that which he quotes as a theorem – that his colleague must either accept both or reject both – is actually flawed, and flawed in exactly the same way as the companion proof of the principle offered by Tversky and Bar-Hillel, who take Samuelson's side in the matter.<sup>11</sup> Both arguments assume that the series of wagers allow for a stopping option – that after any trial  $n$  ( $<99$ ), the colleague may elect to stop playing. Samuelson proceeds by backward induction: after 99 trials, the colleague will face a single decision to accept or reject; out of consistency with his present preference, he must of course reject; after 98 trials, the colleague will also face a single decision to accept or reject, since he knows he will reject the wager if he should complete 99 trials; so he must reject at 98 as well; and so on. Tversky and Bar-Hillel proceed axiomatically, but also with the same stopping-option assumption. But this is an unnatural way to understand the colleague's proposal. And in any case nothing prevents us considering this possibility – the possibility that the colleague can make an irrevocable commitment to play the full 100 trials of the game *ab initio* – as the contrasting option to the single trial. In other words, SC might be proposing instead to take 100 such bets, as a single bundle – without a stopping option. SC might be proposing to *aggregate* his decisions (to bracket his decisions broadly rather than narrowly, in the approving terminology articulated by Read et al. [1999]). Neither Samuelson nor Tversky and Bar-Hillel consider that proposal.

So what, exactly, can be said in support of Samuelson's contention that SC can accept such a bundle if and only if he is willing to take each bet in the bundle individually? It is not entirely obvious. Samuelson's contention was that his colleague's pair of judgments is inconsistent (at least with EU): a person who turns down a particular one-off gamble should also turn down an offer of a string of such gambles. "Each outcome must have its utility reckoned at the appropriate probability; and when this is done it will be found that no sequence is acceptable if each of its single plays is not acceptable. This is a basic theorem." To keep on the right side of EU theory, he thought, one must have approximately the same risk attitude towards an aggregation of independent, identical gambles, as towards each of the gambles taken singly. But how exactly is this principle to be supported? It does not follow directly from EU, since EU is concerned only with single

---

<sup>11</sup> To be clear, Samuelson's proof is different proof from that of Tversky and Bar-Hillel. It shows that bets which are unfair in the utility metric cannot become fair by being compounded. I am simply emphasizing the commonalities between their respective proofs.



choices from among a slate of options, and not with aggregates of choices. Apart from very thin consistency constraints over *basic preferences*, EU (to its shame, as this essay hopes to prove) really has nothing to say about how *choices* must be aggregated, even as it takes a stand on how attitudes towards prospects must depend on (nonoptional) attitudes toward the component outcomes.

One candidate principle to support Samuelson's contention is what I will refer to as the *simple bundling principle for aggregate gambles*:

$$E(A_1 \wedge A_2) = E(A_1) + E(A_2).$$

But why should one accept this principle?<sup>12</sup> One can imagine Samuelson stalwartly defending the simple bundling principle as follows: The single wager belongs to a certain class of such wagers; and whatsoever properties are possessed by the class, as a class or population – such as for instance its statistics – would be inherited by its members. Since the simple bundling principle follows directly from treating the aggregate as a sequence of gambles, each with the properties of its class, or simply as a single (but bundled or class-like) member of that class, the simple bundling principle must be correct.

Applying this analysis to a slightly simpler example than the 100 tosses proposed by SC, consider the example of just two tosses. Suppose that Samuelson offers to allow his colleague to play twice. If SC accepts, then he can anticipate the outcome being one of four: HH, TH, HT and TT (with H standing for heads and T for tails – with equal probability, given the coin's fairness – namely,  $P=1/4$ ). The generator of this distribution is still a fair coin mechanism, a *gamble* which we can represent this way:

---

<sup>12</sup> The problems with this simple bundling principle go way beyond the simple technical one – namely that the principle makes the implicitly assumption that the utility function is normalized such that  $u(0) = 0$ , which is usually not true (since in numericalizing we care only about changes or (ratio) differences in utility, not its nominal value). The body of the paper is concerned with philosophical issues that would arise even if this technical problem were solved.

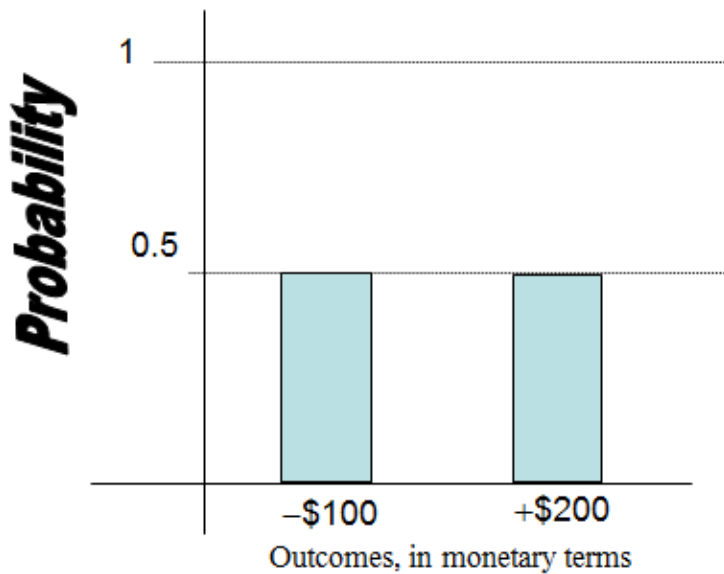


Figure 1. The outcome Heads results in a win of \$200 (the outcome on the right), with relative frequency  $\frac{1}{2}$ , and the outcome Tails results in a loss of \$100 with the same relative frequency

The expectation of the two-fold sequence is easy to compute:  $\$400 \times (\frac{1}{4}) + \$100 \times (\frac{1}{4}) + (-\$200) \times (\frac{1}{4}) = \$100$ . It is robustly positive, whereas not accepting the offer has an expectation of exactly zero, so the former is preferred by EU if we consider only monetary gains and losses (remember we're not at this point taking account of aversion to risk). But notice also: expectations do in fact obey the simple bundling principle, since the  $EU(1 \text{ toss}) = \$50$ . We can further confirm the principle by examining sequences with higher numbers of tosses. The number of potential outcomes grows geometrically with the number of proposed tosses, so the computations get more cumbersome, but straightforward:  $E(3 \text{ tosses}) = \$150$ ;  $E(4 \text{ tosses}) = \$200$ ; and so on.

The simple bundling principle acquires its halo from examining cases like this one, where the statistics for the outcomes in the event class on which gambles are proposed is *normally distributed* – that is, where the distribution function resembles in shape the distribution in Figure 2:<sup>13</sup>

<sup>13</sup> We get in trouble when we apply assumptions like this to contexts where the distribution is manifestly *not* normal, such as in the St. Petersburg game. See Thalos, Richardson (2013) for further discussion of the inadmissibility of using normal statistics in extra-normal circumstances.




Figure 2. The binomial distribution generated by (for example) tossing a fair coin 100 times. The horizontal axis represents the *random variable* being measured—here, the number of Heads in a given 100-toss experiment, and the vertical axis the probability or frequency of that number of Heads appearing in a given 100-toss experiment (for example: in  $\frac{1}{2}$  of all cases at  $N=100$ )

Now the simple aggregation principle can be supported—as it seems Samuelson wished to support it—by a suitable construal of how individual gambles inherit properties from the statistical classes in which they fall—and in particular, how they inherit the properties of expectations. As we’ve seen, the expectation of return *per toss* remains constant, whether we are considering just one toss or 100, or anything in between, as long as we assume the simple aggregation principle. And hence the simple aggregation principle can be invoked in support of the idea that someone prepared to take a set of 100 identical wagers, can turn down any *single* such wager only on pain of a flaw in reasoning about valuation—at least at relatively constant wealth levels. (This by itself should strike the reader as absurd: if an agent is willing to accept a given wager—say, to purchase a lottery ticket—the agent must, on pain of inconsistency, continue to accept such offers as long as they are available at the original price, until such time as the agent’s wealth level or her totality of disposable income changes substantially.)

In one regard Samuelson is right: if one construes the aggregate bet simply as a string of identical decisions (namely, a string of decisions like the one Samuelson offered his colleague), with identical characteristics, then obviously the same considerations must be played out in each situation, leading to identical solutions so long as the context does not change substantially from one situation to the next. And to reiterate Samuelson: “when this is done it will be found that no sequence is acceptable if each of its single plays is not acceptable.” But once one concludes that one must solve each problem the same way, then one must—mustn’t one?—treat the decision about the aggregate as a simple aggregate of single decisions: accept all or reject all. It would be entirely unreasonable to consider some of the offers as acceptable and others unacceptable: they are all the same.

But must one treat a bundled set of bets with the same statistical characteristics simply as a string of identical bets with those statistics? Or can a bundle – as a bundle – deserve a different analysis? This is the question at the core of our concerns in this essay. It shall be my contention that the principle Samuelson enunciates is mistaken: “no sequence is acceptable if each of its single plays is not acceptable” is an invalid principle. This is partly because the simple bundling principle is not generally true,<sup>14</sup> but partly for other reasons as well – for instance, that the agent is restricted, through some constrictions or other, to playing the game a very few times (or even just once). We shall focus here on the latter reasons for the failure of the bundling principle, where focusing on the statistics of an indefinitely large set of identical gambles is simply misguided (and expectational calculations pointless). I shall argue that some sequences are acceptable, even if none of the single plays are individually acceptable. And perhaps less contentiously, some sequences are unacceptable despite each of its individual plays being individually acceptable: it is surpassingly rational, for instance, to stop buying lottery tickets after one has already purchased 100, even if no one single lottery ticket purchase is individually unacceptable and even if one has the disposable income to do it. Enough is enough. Such network effects emerge among independent events whose aggregate effect on the valuation of the prospect is a nonstraightforward (more precisely, nonlinear) function of the individual events.

The fundamental and ultimate moral shall be this: decision points are networked. Not surprising since prospects themselves are networked: the value of one condition or outcome to an agent has to be made in the context of all other conditions and outcomes enjoyed by the relevant agent over all time. And so we must proceed with making decisions accordingly. A very simple example from consumer economics: the value of a given car tire to an agent depends on how many other tires the agent has already acquired, and whether they are associated with an automobile that the agent has also acquired. But we would like to see how this applies very generally. And as with all networks, one must watch for network effects – and they are not always predictable from present local conditions. A global perspective is required to make sense of how valuation-relevant considerations will aggregate. It is my contention that risk-numericalizing theories of valuation have not risen to this challenge.<sup>15</sup>

---

<sup>14</sup> We shall not examine such cases here, but leave them for another occasion.

<sup>15</sup> It is my own conviction that no purely formal theory can do so, but that is an argument for another time.

While examples are useful for generating doubt, we still need to make the principles at work in the examples I have invoked for illustration crystal clear. And to see precisely where they direct us. The remainder of this essay is devoted to these tasks.

#### 4. False expectations?

I shall argue in this section, and using very simple gambles, that the monetary value of a gamble involving only money is not always expectational, and so not related in the simple way Samuelson envisaged to the statistics associated with the class in which it falls. This analysis will provide a stepping stone to more nuanced analysis of the differences between expectational reasoning and instrumental reasoning.

First, we will address the following question: when is utilizing an expectational calculation *most* defensible? The obvious answer is: when the expectational formula matches exactly the outcome of the action. This can occur in one of two ways: (1) when the outcome is certain (no probabilities intermediate between 0 and 1 are involved); or (2) when the mix of outcomes one actually suffers or enjoys is exactly like the mix of outcomes anticipated by an expectational calculation. (1) is self-explanatory. (2) concerns the *faithfulness* of a given sample as a statistical representative of the entire (potentially infinite) population. To appreciate the stringency of the faithfulness condition, one needs to think in terms of populations of outcomes, the sorts of populations that statisticians discuss under the umbrella term *distributions*.

Statisticians now characterize statistical distributions—more formally, probability distribution functions (PDFs)—using a handful of parameters called *moments* that are collectively meant to give a measure of the shape of a scattering of points or functionals like that in Figure 2, as well as more messy measurement or observational data. The one in the binomial distribution of Figure 2 is well-estimated by the most familiar of bell-shaped curves—the Gaussian (the so-called “normal”) distribution, poster child for statistics itself.

The first moment of a distribution is its mean or weighted average—the number associated with the expectation of EU theory. In the distribution of Figure 1, the mean is 50—that is,  $N \times P(\text{Heads})=100 \times \frac{1}{2}$ , where  $N$  is the number of trials (or bets in this case). (This mean is also the largest-probability event, so in this instance it is also the *median*.) The second moment—and most important for our purposes here—provides a measure of the *width* or *spread* of the set of points of the distribution around its mean, in this example it would be the quantity known as the *variance*. The variance of our binomial distribution of Figure 1 is  $N \times$

$P(\text{Heads})[1 - P(\text{Heads})] = (\frac{1}{4})N=25$ . (A third moment, that won't demand attention here, provides a measure of the *skewness* – the extent to which the distribution favors one or the other side of the mean. There are a handful of further moments that we won't be concerned with here.)

The first point is that distributions can differ along a number of different and independent dimensions—along different *moments*. Importantly for us here, two distributions sharing the same mean, can differ in their variances. Consider for instance the two distributions generated from the same-mean gambles (Left and Right) depicted in Figure 3. The relevant difference between the two is that Left gives players very small chances at each of \$1 and \$9, and comparatively larger chances at nothing, while Right gives players roughly even chances at the comparatively much smaller sums of \$0.6 and \$1.6 respectively. To reiterate, the means (expectations) are equal. But decidedly not the variance. One can see by simple inspection that the spread in Left is much larger than that in Right. (And using  $N=100$  in each case, Left has a variance of 82 while Right has a variance of 24—a near four-fold difference.) What is the significance of this difference?




Figure 3. Two games/gambles (Left and Right) with the same mean ( $\mu = 1$ ) but different variance: Left has variance of 82, while Right has a variance of only 24, using  $N=100$

Suppose that it costs \$0.5 to play either game of Figure 3—to enter into either of the two gambles for the small price of \$0.5. The means are identical, at \$0.5 each—either gamble will return \$1 for every half-dollar staked. Consequently, EU commends indifference between the two gambles (we're still assuming that money is the only thing at stake, and that utility is a linear function of money). Imagine now that you have \$500 to stake—that you have budgeted that sum for gambling—then you might quite rightly be indifferent between the two gambles, figuring that if it costs \$0.5 to play, then you can expect to double your money by playing all \$500 on either game: you'll end up with  $(100 \times 1 + 100 \times 9)$  dollars by going Left and

( $600 \times 0.6 + 400 \times 1.6$ ) dollars by going Right – \$1000 either way – with rather high probability *because you will be playing either game 1000 times*. But if instead the game were more costly as well as more profitable by a hundred-fold, as depicted in Figure 4, and you still have only \$500 budgeted to the game, then you might well figure: if it costs \$50 to play, then by playing only ten time, you are quite a bit more likely to end up with absolutely nothing (a loss of your entire stakes) if you play Left, whereas playing Right will at worst yield \$600 (for an overall gain of \$100), so that Right is much to be preferred, for reasons that are entirely to do with money, in spite of the equal expectations.

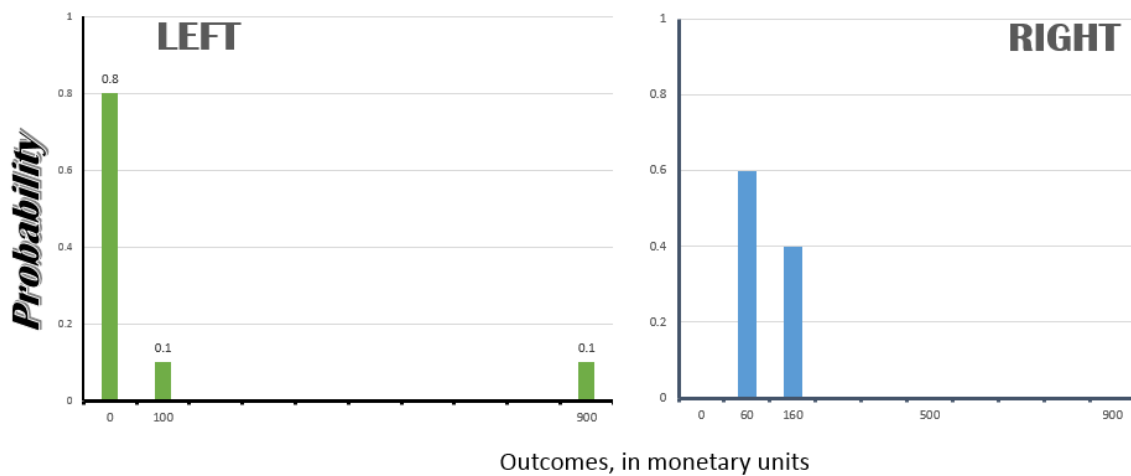


Figure 4. Another set of gambles, like those of Figure 3 but larger in scale by 100-fold

This example demonstrates decisively that the statistics of a gamble can matter *over and above its contribution to expectations*, even when the only thing at stake is money. Let's now see what happens when we take our analysis to its logical conclusion: what if the gambles available are those in Figure 5, rather than those of Figs 3 or 4?




Figure 5. Another permutation on the gambles of Figure 3, scaled by 1000

Now it costs \$500 to play, so you can only play once. Now it doesn't make any sense to say that one can "expect" a return of \$1000 on every \$500 you gamble, although this is still what the expectation principle says. When you are choosing between Left and Right in Figure 5, there are only 5 outcomes, one of which is \$1000 (with probability 0.1 if you go Left). But if you go Left, your most likely outcome is loss of your entire stakes. Whereas by going right you assure yourself a gain of \$100, but might (with probability 0.4) come away with an additional \$1100 more than you staked. All you care about here is money. What do you do?

In each of Figures 3-5, established terminology has it that Left is a *mean-preserving spread* of Right. Suppose you choose Right. Then you are choosing the (*second-order*) *stochastically dominant* option, and you, dear reader, are *risk-averse*. It's a matter of your taste for risk. According to prevailing wisdom on these matters, someone with opposite "inclinations" to yours is no more and no less rational. De gustibus non disputandum.

It is not my contention here to challenge the prevailing wisdom on rationality. I wish instead to focus on *just what sort* of rationality is being exhibited by the person who happily accepts either the wagers of Figure 3, but declines some or all of those in Figures 4 and 5. What sort of reasoning is it that lies behind SC's willingness to accept a bundle of identical wagers when he was not prepared to accept a singleton on its own? My answer, no doubt to the horror of Samuelson, is simple: it is unimpeachable instrumental reasoning.

It should be a principle beyond controversy that the value of some option to an agent has to be gauged in relation to something else: to the goals or aims of that agent in selecting that option. There is hardly any principle more bedrock in the area of instrumental reasoning. Now, it might well seem as though we are catering for this bedrock principle when we are computing expectations. But we shall argue in this section that this is simply not so.

The assumption that valuations must be expectational, as might now have become clear already, rests on an assumption about the relationship of the agent to the potential outcomes ( $O_1, \dots, O_n$ )—namely that the agent is in a position to view each such outcome as making a nonnegligible contribution to his or her utility "purse"—a contribution in proportion to its probability of being realized. This principle, however, rests on the idea that every possible outcome makes an appropriately weighted contribution to the value of the gamble as a whole. *But this principle is valid only if one is expecting to play "long enough" to realize the probabilities in question in their true proportions.*

Samuelson (1963) saw clearly the true relationship between expectational valuation and repeated play when he wrote: "Each outcome must have its utility



reckoned at the appropriate probability; and when this is done it will be found that no sequence is acceptable if each of its single plays is not acceptable.” This is the idea that expectational reasoning must be premised on the condition that the agent is realizing all the probabilities in their (true) proportions. But now this raises an important question: why should someone who will be realizing only one out of the numerous possible outcomes, because s/he is in a position to play only once (or, more generally, a relatively small number of times), have to “reckon all the utilities at their appropriate probabilities”? They should not be so required, or so I will contend. I claim that in fact focus on probabilities exclusively can distort our appreciation of what someone in a context of risk needs to take into account. More precisely, exclusive focus on the “relevant probability” displaces other important considerations. The series of gambles in Figures 3-5 above demonstrate concretely that statistics can matter beyond their contribution to expectations, even when all the agent cares about is money. But what is contributed over and above a calculation of expectations? This is the question to which we now turn.

Someone who is willing to take either Left or Right in the game in Figure 3, because they know they can play 1000 or more times, might well turn down *Left* in both Figures 4 and 5, because they don’t have the resources to *utilize those games as a means of growing their stakes, while nonetheless accepting Right in both*. This is not a matter of mere taste for risk. It’s a matter of *capitalization vis-à-vis their goals*.<sup>16</sup> If you decline the Left gamble in Figure 4, you may simply be realizing that you’re *insufficiently capitalized* for doubling your money, given the tools you have available to you to do it. Whereas you could indeed aspire to double your bankroll with either gamble of Figure 3.

This is the point: either of the gambles (Left or Right) in Figure 3 constitutes a *good means* to your ends, a *good instrument* of growing your cash. But the same cannot be said of all the games in Figures 4 or 5 – *Right* might be good means to that same end, but *Left* not so much. The gamble of Figure 5 Left is an especially poor instrument. It is a gamble in the colloquial sense of the term (as well as the technical sense), and that fact stands as a barrier to its being an instrument to your goals (in every sense of the term). The term *instrument* is – if anything is – a contrary to that of *good luck*.

Let us now return to Samuelson’s colleague SC. He is confronting the gamble of Figure 1, and (apparently) is prepared to stake the \$10,000 it takes to play the game 100 times. At  $N=50/100$  wins (which is the mean of his gamble), his total payout is \$20,000. He can reasonably view this game as an opportunity to double

---

<sup>16</sup> The terminology derives from Thalos and Richardson (2013).

his cash. But if he can only play the game once, he might very well view this as a poor means to making money. But if he can choose to play either once or 100 times, and if his aim is indeed to make money (as we can suppose for purposes of argument that it is), is he not more reasonable to decline the one-off while accepting the 100?

It is clear why Samuelson's error *vis-à-vis* his colleague's reasoning is so monumental: his argument purports that there can be no relevant differences between long runs and short runs, so long as we stay roughly within a certain range of wealth. This is an enormous mistake, as considerations of statistics show the importance between short and long runs, and show furthermore that these considerations are independent of levels of wealth, but instead concern the agents' specific goals in view. And this is in no way a mere matter of taste. Expectational reasoning, when it tells us that the difference between same-mean gambles is a matter of taste, stops short; it stops before it reaches full-fledged instrumental reasoning. Fully realized instrumental reasoning, by contrast, takes into consideration what I called the *faithfulness* condition (the resemblance relation between full distributions and sample populations drawn from full distributions) to take the measure of an option as an *instrument* toward fulfillment of goals. We will turn to the subject of how to distinguish between instruments and noninstruments in the next section.

## 5. Putting instruments back into instrumental reasoning

Certain bargains are categorically unattractive—to take a concrete example type, consider gambles involving the lives of one's children. These are unattractive because (among other things) one cannot with any seriousness consider taking repeated gambles with lives; for once the gamble is lost just once, there are no further gambles to take—no way to come out ahead in the end. One must consider gambles with one's children's lives (where one is obliged to take them, for example by dictates of custom or culture in places where female genital mutilation is an absolutely everyday event) as inherently single-sequence gambles. And this explains (at least in part) why these sorts of gambles are—and indeed should be—inherently unattractive.

Such gambles cannot be employed as instruments to come out ahead in the end. By some contrast, buying lottery tickets is not a *good* means to wealth, but it should not to be set aside entirely as an instrument. The aim of this section is to explain these points utilizing the vocabulary of statistics. Instrumental reasoning requires us to think in these terms; by contrast, expectational reasoning demands that we do without them.

A hammer is not a good means to the end of driving home a screw. Even where it succeeds, much damage will be done. But more importantly, using the hammer runs a much greater risk of spoiling the work entirely – everything you aim to accomplish in the enterprise. But if it’s all you have at hand, and the screw has to be driven or else the world collapses, you might use it anyway, hoping for good luck. For in case you succeed, it will most certainly be entirely by good luck, and not because you used a proper instrument. These distinctions are important in practical reasoning.

Why is the hammer no instrument at all, whereas a poor screwdriver is merely a poor instrument? For this reason: the hammer, in *addition* to giving a poor probability of success, *also* runs you a fair to sizeable probability of absolute failure in relation precisely to the goal you are trying to reach; it courts disaster as well as moderate success. And harm is done either way. We can display this point visually as follows (Figure 6):




Figure 6. Specs for a hammer (left) and screwdrivers (right). A is the outcome “expected” from a high-quality screwdriver, B, C and C- from lesser screwdrivers. Z is disaster. Notice the relative likelihood of disaster on the left. Moreover, the variance on the right will always be considerably smaller

For simplicity, I’m assuming binomial distributions. The “specs” of Figure 6 help us discern the instruments (whose specs appear on the right) from the noninstruments (whose specs appear on the left). Binomial distributions of noninstruments are characteristically many times those of instruments for production of the outcomes denoted here by A-C-. Noninstruments court disaster (Z), whilst contributing to the wide variance. And if an instrument courts disaster at all, its contribution to disaster is small to negligible. (The lottery ticket, as a means to wealth, rates a D or F, but it is no “hammer” unless the customer has nothing but the purchase price of the ticket to her name.)

Not coincidentally, cases of risking one's children's lives (for instance, via female genital mutilation) are easily characterizable via Figure 6 Left, if we interpret the horizontal axis as representing the would-be victim's well-being.

So by looking at the PDF of a gamble, and using statistical concepts, we can render more nuanced appraisals of whether the gamble in question provides us with good "specs" *vis-à-vis* the job we'd like to use it for. And risk-numericalized reasoning, because it ignores so much of the information contained in a full PDF, misses out a great deal of instrumental reasoning. Calculating an expectation is just the beginning (or middle, if you prefer) of instrumental reasoning; it's not the end.

## 6. Fidelity: The difference N makes

And we mustn't stop at "specs". There's still more that statistical concepts can give us. This is the concept of *fidelity* itself. We have seen that N—the number of times one exposes oneself to a given risk—matters a great deal. Expectational reasoning misses out this aspect of instrumental reasoning. Paying attention to N allows one to take into consideration the extent to which one's own case will bear resemblance or fidelity to the statistical distribution generating the probabilities that enter into the calculated expectations. It allows one to take a truer measure of one's situation as generating a faithful sample of the relevant statistical population that constitutes the risk one is taking.

What one is doing, in any given gamble, results directly (at least in part) from how many times the gamble is, can be or indeed must perforce be taken. In general, an agent must attend to how the present decision is networked with other decisions, to constitute an instrument in service of the agent's goal. This is an important feature of reasoning instrumentally—it is part of what actors must attend to in service of their ends, since it quite frequently takes more than one action to achieve a goal—keeping your teeth healthy takes multiple and relevantly spaced episodes of maintenance—and these must be coordinated. To provide genuinely useful assessments of gambles, we must be sensitive to coordinations that must be made amongst actions, in service of larger goals.

## 7. Objections and replies

One response at this point might be that giving attention to anything besides the mean (the first moment) of a statistical distribution involves allocating concern or regard for goals *outside* of those that are explicitly defined in one's preferences—one's utility function. For instance, giving attention to variance or skew might involve considerations for *caution*, *safety* or *security*. But not all agents

care about these things to the same degree, and the same agent might not care about these things in the same way in all contexts. There is no reason to give attention to these things if the stakeholders in the situation do not.

This is fair enough. And indeed please note that I am not suggesting that stakeholders *must* give attention to such things, but only that they may.<sup>17</sup> But even granting the point that paying attention to such things suggests there are goals, perhaps overlooked in computing utility, to which a stakeholder is (now) paying attention, we should nonetheless demand a very good answer to the question of why the mean (the statistical average) deserves a special status, a place of honor in instrumental reasoning – if we insist upon its use to the exclusion of other considerations. I’ve argued above that variance and N deserve consideration in instrumental reasoning. Why should they be denied a place of honor? The advocate of exclusive expectational reasoning cannot absent herself from engaging at this point. We require an answer as to why risk-numeralization à la EU absolves agents from having (or getting) to consider other features of their situation as well.

This is the challenge recognized squarely by Maurice Allais, who was among the very first critics of EU. Allais was always mindful of the power of context, especially the risk landscape, in appraisal of risky prospects. (And Allais did not subsequently conclude – as so many decision theorists hastily do today – that people are irrational in their dealings with risk. Among the most egregious is Ariely.<sup>18</sup>) Instead Allais objected to EU on grounds that it failed to take risk itself seriously, as an independent psychological “consumer good” with its own independent economy, writing in his Nobel Prize lecture that EU “struck me as being unacceptable because *it amounts to neglecting the probability distribution of psychological values around their mean*, which precisely represents the fundamental psychological element of the theory of risk.”<sup>19</sup> EU did not match Allais’ preferred conception of valuation – more precisely, it did not capture his conception of the way that caution infuses all aspects of deciding in multiple ways. EU neglects a “very profound psychological reality, *the preference for security in the neighbourhood of certainty*.”<sup>20</sup>

---

<sup>17</sup> By contrast, the approach known as “mean-variance” first advanced by Markowitz (1952) and developed by Kroll et al. (1984) and again by Blavatsky (2010) requires attention to certain statistical moments of a distribution (the first and second). And to repeat: such attentions are not sufficient to account for the effect of N.

<sup>18</sup> Ariely (2008, 2010).

<sup>19</sup> Allais (1998).

<sup>20</sup> Jean Hampton (1994): 235 makes a related complaint, *vis-à-vis* EU’s treatment of preferences. According to Hampton, EU is a theory divided against itself. She writes: “While it [EU] is supposed to permit our preferences over actions to be partly a function of attitudes toward risk, one of its

Security, Allais seemed to think, matters to at least some people, independently of other things that might also matter to them. Safety of that which one cares about is itself a value, he reckoned, over and above the value of the object whose safety is threatened. When it is, EU is a poor approximation of instrumental reasoning – at best. And who is to criticize people with such concerns? After all, and following Russell, we are purporting not to pass judgments – not in instrumental reasoning, anyway – on the goals people choose to adopt and pursue. Allais thought that EU was passing such judgments, and accordingly thought EU is flawed. From his ideas springs an entire branch of theories of choice under risk – meant at least at their point of origin as descriptive – now referred to as non-expected utility theories.<sup>21</sup>

While I have sympathies with that school's founding thought (to the effect that EU might not be the only admissible game in town), there persists the question whether there is not a larger principle at stake. It is my conviction that there can be more than attitude to risk at stake in the refusal to adopt a risk-numericalized calculus as the measure of instrumental reasoning – larger things still.

## 8. Planning

To appraise a bargain, one needs to do more than compute its expectation value: one needs also to appraise the gamble as a means to attaining the goals in view, and this assessment must contrast the other options for achieving those goals. And there is more: one also must take stock of how the present decision will be intersecting with future decisions, in service of these goals. One must, in other words, be mindful of one's plans.

Someone will be sure to remind us that plans too (and in general decision trajectories over time) can be subjected to risk-numericalized utility computations. And that all risk-numericalizing theories of decision prescribe the application of this calculus to plans as well as to individual decisions.

---

central axioms requires that our preferences over actions be *purely* a function of our preferences over consequences." Hampton's point is that, whereas EU is prepared to render judgments utilizing any utilities you please to specify, whether you care about consequences only, or care also about the states in which you enjoy them, it nonetheless demands that your utilities for outcomes be state-independent at *some* level. According to Hampton, EU requests that, at some level, your preferences be purely over the outcomes, rather than also over the conditions in which the outcomes come out.

<sup>21</sup> Machina (2004); Starmer (2000).

But there are numerous philosophical questions that are not addressed by the principles of risk-numericalizing reasoning. Here are a few:<sup>22</sup>

1. Should decisions always be taken sequentially (one at a time), or can they be bundled in any way? If the latter, what rules govern bundling? Is bundling always permitted, or is it sometimes disallowed?
2. Are decisions made at a given time ever subject to challenge at a future time? If so, are decisions *always* to be challenged? What rules govern the process of deciding to challenge decisions made at an earlier time?
3. Is it possible to bundle together decisions of different parties? Or is it never permitted? What rules govern such bundling?

All these questions pertain to the simple principle of bundling introduced earlier:

$$E(A_1 \wedge A_2) = E(A_1) + E(A_2).$$

Allow  $A_1$  to denote adoption of the overall plan, and  $A_2$  to denote an element of the plan to be performed at a date future to adoption of the plan. Allow that  $A_1$  and  $A_2$  might involve different actors. One can now ask how one ought to decide about carrying out a part of a plan already adopted. And no risk-numericalizing theory of decision answers the question. This is because, as we noted before, all existing theories of risk-numericalizing decision are single-decision theories: they have nothing whatever to say about how decisions ought to be bundled, and so nothing whatever to say about how decisions should be connected over time. They do not pronounce on the simple bundling principle, neither in favor nor against. They are theories of serially-independent decisions. And so they are stunted, as theories of planning. If risk-numericalizing calculi are to have any hope of providing true service in planning contexts, they will have to address such questions.

It might be suggested at this point that theories under the umbrella of “game theory” might come to the rescue of EU. After all, such theories are meant to take into effect the dynamics of interactions, and social interactions in particular. I respond, however, that game theoretical analyses routinely take a single-player decision theory (namely, EU) for granted, and build upon it principles of multi-player interaction. In other words, what game theories try to do is take into account the strategic reasoning involved in making decisions in social and other multi-player settings. But this is nowise to take into account the networking between decisions that we have gestured at – the network between decisions *of a sin-*

---

<sup>22</sup> Many of these questions have already come to philosophical awareness in the form of paradoxes or puzzles, such as for instance Kavka’s (1988) Toxin problem. See Thalos (2008) for exposition of problems *vis-à-vis* planning and their relations to certain philosophical puzzles.

gle decision maker over a lifetime. Now, it might be suggested at this stage of argument that a game theoretic analysis can take this into account if it treated the single player as a series of separate players whose existences spanned only a small interval of time. This is a promising strategy – it is the strategy routinely utilized by researchers aiming to understand the psychology of real-life decision makers in the wake of Ainslie’s observations.<sup>23</sup> It seems to promise that we can produce an analysis of the networked quality of a single decision maker’s decisions over a lifetime, as a function of individual decisions made sequentially in the usual way. Unfortunately, this strategy also requires that the concerns of each agent associated with an individual decision (the concerns of any given time-slice agent) are *not* focused through a lens of interconnections: each individual in the sequence that makes up a four-dimensional agential entity, is treated as someone whose individual decisions are not networked. But we have argued above that treating individual decisions as not networked is precisely the problem. Why should we think that multiplying this assumption, and spreading an individual over time in this artificial way, can be a means to its solution? (Indeed, I would maintain that this analysis prompts the befuddlements reported in Ainslie in the first place.) I cannot think of a reason for supposing it would do anything but multiply the error. It will perhaps help to take a specific example. SC’s case has served us well so far; and it will serve again at this juncture.

SC declines – at a single moment in time – to take Samuelson’s first offer, and counteroffers with one of his own. I cannot imagine how SC’s decision can be construed as eligible of analysis on the sort of four-dimensional analysis being proposed here. For it would have to treat SC as at least two time-separated agents. And this just doesn’t even begin to get purchase on the problem. Still, there might be a way forward with the proposal that I have not been able to grasp. So I leave it here as an exercise for my potential opponents to work out the details of how to rationalize SC’s decision in a four-dimensional framework.

## 9. Conclusion

Expectational reasoning is not full-throated instrumental reasoning. This is by no means to say that it is rubbish. One can view an expectational calculation as a *first moment* in instrumental reasoning. But there are numerous cases where one cannot stop with a first moment.

The moral of this essay is really to exhibit the richness of instrumental reasoning – to display the fact that we have hardly scratched its surface when we

---

<sup>23</sup> Ainslie (1975).



have produced an expectational calculation, whether after a classical fashion or after the fashion of more nuanced non-expectational calculi. Risk-numericalized calculations are not always in harmony with considered human judgment – as now a very large literature in behavioral economics will attest. And while many behavioral economists move quickly to the conclusion that human judgment is flawed – indeed, positively perverse – we should not rule out that such human judgment might well be better attuned to the demands of instrumental reasoning than are the risk-numericalizing theories of decision.<sup>24</sup>

## References

- Ainslie G. (1975), “Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control,” *Psychological Bulletin* 82 (4): 463–496.
- Allais M. (1998), “An outline of my main contributions to economic science,” paper presented at the Nobel Lecture, Ecole Nationale Supérieure des Mines de Paris et Centre National de la Recherche Scientifique, URL = [https://www.nobelprize.org/nobel\\_prizes/economic-sciences/laureates/1988/allais-lecture.html](https://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/1988/allais-lecture.html) [Accessed 5.9.2017].
- Ariely D. (2008), *Predictably Irrational: The Hidden Forces That Shape Our Decisions*, Harper Books, New York.
- Ariely D. (2010), *The Upside of Irrationality: The Unexpected Benefits of Defying Logic*, Harper Books, New York.
- Blavatsky P. (2010), “Modifying the Mean-Variance Approach to Avoid Violations of Stochastic Dominance,” *Management Science* 56 (11): 2050–2057.
- Buchak L. (2014), *Risk and Rationality*, Oxford University Press, Oxford.
- Frederick S., Loewenstein G., Donoghue T. (2002), “Time Discounting and Time Preference: A Critical Review,” *Journal of Economic Literature* 40: 351–401.
- Hampton J. (1994), “The Failure of Expected-Utility Theory as a Theory of Reason,” *Economics and Philosophy* 10 (2): 195–242.
- Kavka G. (1983), “The Toxin Puzzle,” *Analysis* 43 (1): 33–36.
- Kroll Y., Levy H., Markowitz H.M. (1984), “Mean-Variance Versus Direct Utility Maximization,” *Journal of Finance* 39 (1): 47–61.

---

<sup>24</sup> This paper constitutes a follow-up to investigations of the shortcomings of EU as the sole decision principle in contexts of risk, begun in Thalos and Richardson (2013). I would like to thank my co-author Oliver Richardson for beginning the conversation. Thanks to Chrisoula Andreou, Tyler DesRoches, Don Ross and Jonah Schupbach for rich conversations about topics surrounding risk. Tomasz Żuradzki and a number of anonymous referees for this journal provided invaluable suggestions and corrections. For all the errors and other missteps that may remain, despite the wealth of assistance offered, I have only myself to thank. Research support was provided in part by the National Science Foundation, USA, under grant number SES-0957108, and by a grant no. 2015/17/B/HS1/02279 funded by the National Science Centre, Poland.

- Machina M. (1982), "'Expected Utility' Analysis without the Independence Axiom," *Econometrica* 50 (2): 277-323.
- Machina M. (2004), "Nonexpected Utility Theory," [in:] *Encyclopedia Of Actuarial Science*, vol. 2, J.L. Teugels, B. Sundt (eds), John Wiley & Sons, Chichester: 1173-1179.
- Markowitz H. (1952), "Portfolio Selection," *Journal of Finance* 7 (1): 77-91.
- McClennen E. (1990), *Rationality and Dynamic Choice: Foundational Explorations*, Cambridge University Press, Cambridge.
- Quiggin J. (1982), "A Theory of Anticipated Utility," *Journal of Economic Behavior and Organization* 3 (4): 323-343.
- Pettigrew R. (2015), "Risk, Rationality and Expected Utility Theory," *Canadian Journal of Philosophy* 45 (5-6): 798-826.
- Read D., Loewenstein G., Rabin M. et al. (1999), "Choice Bracketing," *Journal of Risk and Uncertainty* 19 (1-3): 171-197.
- Samuleson P. (1963), "Risk and Uncertainty: A Fallacy of Large Numbers," *Scientia* 98: 108-113.
- Starmer C. (2000), "Developments in Non-expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk," *Journal of Economic Literature* 38 (2): 332-382.
- Thalos M. (2008), "On Planning: Toward a Natural History of Goal Attainment," *Philosophical Papers* 37 (2): 289-317.
- Thalos M., Richardson O. (2013), "Capitalization in the St. Petersburg Game: Why Statistical Distributions Matter," *Politics, Philosophy & Economics* 13 (3): 292-313.
- Tversky A., Bar-Hillel M. (1983), "Risk: The Long and the Short," *Journal of Experimental Psychology: Human Learning, Memory, and Cognition* 9 (4): 713-717.