Combining portfolios of risky research projects within a science of science funding

Luigi Buzzacchi
DIST and FULL
Politecnico di Torino, c.so Duca degli Abruzzi 24, 10129 Turin, Italy
luigi.buzzacchi@polito.it

March 2022

Contents

Abstract

In recent years, scholars and policy makers have been questioning the efficiency of the allocation of funds granted by governmental agencies and private foundations to scientific research programs. The selection procedures used worldwide by most of the programs remain similarly and uniformly organized despite the various critiques raised in scientific and public debate. Two main concerns centre on the poor treatment of uncertainty in the selection process: on one hand, the evaluation of projects is mainly focused on 'expected' quality, with a lack of analysis on their riskiness; on the other hand, a 'one-to-one' approach prevails, i.e. projects are evaluated as if they were selected in a mutually exclusive fashion. Yet, funding agencies actually prove to be risk-averse; moreover, they select *portfolios* of research projects. Consequently, disregarding the risk associated with alternative projects, particularly their marginal contribution to the risk of the entire program, naturally leads to sub-optimal choices, especially considering that individual risks are generally not additive.

(Financial) portfolio theories provide useful hints for portfolio composition, but financial investments obviously show many distinctive features, both technical and institutional, when compared to investments in scientific research. Accordingly, financial portfolio techniques can be adopted only in specific cases and with several caveats.

This paper reviews the concerns of the literature and illustrates how portfolio theory can be adapted into new procedures for selecting research projects. To this end, the reasonable set of assumptions and parameters needed for improving the decision-making process is discussed. The second part of the paper introduces the specific issue of research funding, in particular the fact that research projects often provide discrete and skewed returns. Several corrections to a 'pure' financial approach are then proposed, with special focus on the benefits of stage financing.

1 Introduction

The crucial relevance of scientific research for economic growth and welfare needs no justification. Because knowledge – and especially knowledge created by scientific research – has properties of public good, the correlated amount of investments in the private sector is expected to be suboptimal (Arrow, 1962). Resorting to research programs sponsored by governmental agencies and private foundations is therefore increasingly relevant; nevertheless, these programs are the subject of growing criticism (since at least Viner et al., 2004) with regard to their allegedly inefficient allocation of available funds.

Grant proposals are usually selected using conventional and undisputed procedures that have been around for decades. As such, focusing special attention on scientific research in order to improve the design of grant allocation mechanisms, thus enabling and sustaining innovation, has recently been advocated by many authors. As Franzoni et al. (2021, p. 29) put it "... *important contribution granting agencies can make would be to support research which builds knowledge on the design of funding programs and reviewing practices related to risky proposals that have the potential of delivering breakthroughs*".

Starting from this general standpoint, this paper is based on two specific premises:

- the research projects selection procedures are mainly based on a '*one-by-one approach*', thus disregarding the opportunities offered by the conscious adoption of *risk diversification* practices; however, the social benefits of proper risk management techniques could be more or less relevant depending on the profile of the distribution of research returns; a specific investigation in this regard is therefore more than appropriate;

- the ability of granting agencies and peer reviewers to fund fundamental and breakthrough research is usually considered to be unsatisfactory. In particular, *high-risk/high-return* (**HRHR**) projects are deemed under-represented in the portfolios of granting agencies; understanding the determinants of this failure can indicate how to design more socially efficient selection procedures.

Granting agencies usually select and manage 'portfolios of projects' whose *aggregated* risk is a (non-linear) function of the risk of every single project. Projects are selected with agency-specific procedures. Some common traits can nevertheless be individuated: panels of peer reviewers analyse and evaluate single proposals, the opinions of the panellists are then aggregated and the portfolios are composed through a '*one-to-one approach*'. As Franzoni et al. (2021, p. 16) describe, "... *agencies generally review proposals one by one, rank them in descending order of overall aggregated score, and then distribute funds according to the score until the budget is exhausted*". In any case, the lack of attention focused on portfolio approaches by granting agencies is usually taken as read.[1]

Based on this premise, I will examine the cost of the 'one by one' approach that evaluates single projects 'individually', disregarding that they are collocated into portfolios, and that the risk borne by granting agencies is the risk of the entire portfolio (where the composition of risks is subadditive, thus obtaining *risk diversification*, which can be greatly beneficial for risk-averse subjects).

The benefits of diversification when selecting a portfolio of assets naturally recalls the methodological contributions by von Neumann and Morgenstern (1953) with their axiomatic treatment of choice under uncertainty, and the pioneering work in modern portfolio theory by Markowitz (1952 and 1959) and Tobin (1958). However, these works are properly structured for analysing the allocation into portfolios of 'usual' financial securities, traded on efficient markets (and with normally distributed returns).

As mentioned, the literature concerning risk diversification in the financing of scientific research projects by granting agencies is fairly limited and entirely qualitative. The analogous issue is much more debated in two other streams of economic literature: i) the composition of VC funds, even if in this case the emphasis is mainly on the pros of risk diversification as opposed to the cons of excessive diversification that come at the cost of managerial de-specialization (Norton and Tenenbaum, 1993) and ii) the optimization of the selection and development of concurrent alternative R&D projects in private firms (Baysinger and Hoskisson, 1989). All these literatures are focused on the institutional and technical specifics of its field of application, but all of it indicates promising directions for the challenging design of funding programs and reviewing practices.

The belief that the HRHR research is under-represented in the portfolios of granting agencies is actually based on qualitative evidence and indirect clues. Taking for granted this hypothesis, the academic discussion usually attributes this bias to *behavioural* or *procedural/tactical* arguments. Sometimes the two reasons are mistakenly confused. With for the former, the risk-tolerance of the

[1] Few explicit mentions of the need to consider how risk aggregates into portfolios are provided: "*At a minimum, panels need to think about correlation among the proposals they are funding* " (Franzoni et al., 2021, p. 25); "[ARPA-E and DARPA] *often balance programs with complementary, and potentially competitive, technologies in order to create a diverse portfolio (Fuchs, 2010)*" (Azoulay et al., 2019, p. 81). In any case, selecting projects in descending score order actually neglects *risk* completely, not just *risk diversification*.

decisionmakers turns out to be considered excessively low with respect to a socially optimal level of risk aversion (Azoulay et al., 2011; Nicholson and Ioannidis, 2012; Linton, 2016; Franzoni et al., 2021); the procedural argument, on the contrary, has been raised more recently (Azoulay et al., 2019; Veugelers and Zachmann, 2020; Franzoni et al., 2021) and suggests that granting agencies overlook the opportunities offered by risk diversification techniques, thus overstating the actual risk borne by granting agencies and acting in an excessively cautious manner.

However, as I will argument below, risk diversification – in the *traditional* setting suggested by financial economic literature – does not mean that the probability of funding high risk projects would be higher (rather, the opposite might hold true). Moreover, in the second part of this paper I will also challenge the 'traditional setting': scientific research risk is quite different from financial risk, even from a technical point of view, and the results of scientific risk aggregation do not necessarily conform to the pattern of financial risk aggregation. Finally, from a dynamic perspective (as opposed to a traditional static one), the beneficial contribution of risk aggregation is driven by slightly different mechanisms, and in this setting risk management is expected to significantly contribute to rebalancing the presence of incremental and breakthrough research in the choices of granting agencies.

The issues addressed above require both institutional and technical care.

On the institutional side, the role and the risk attitude of granting agencies must be understood; they are obviously very different from those of mutual funds, VC funds, or R&D departments. In particular, the former aim to provide socially valuable public goods, while the latter have the explicit mission of producing private economic values. This condition implies that granting agencies, as institutions, are not necessarily risk-averse, and even if they are risk-averse, socially optimal risk diversification is not necessarily requested at the level of the individual granting agency. In this sense, the social cost of suboptimal risk management practices is not obvious.

On the technical side, to the best of my knowledge, the analysis of the effects of risk diversification within a portfolio of scientific research projects is new in the literature. However, no relevant results in this direction can be obtained unless some general statistical properties of projects' returns are made explicit. The common adage is that if the risk is reduced, then investors will be able to finance more risky projects, and the greater the risk, the higher the expected return.[2] This is not necessarily

[2] "*The 'one by one' approach typically used in panels works against selecting risky proposals*" (Franzoni et al., 2021, p. 25) and "*... greater diversification reduces* [VC] *fund risk, enabling risk-averse managers to select riskier investments in the first place and, thus, investments with higher expected returns*" (Buchner et al., 2017, p. 520).

true: in general, it is risk aversion that works against the selection of risky proposals, not the riskiness of single assets; the 'one by one' approach simply neglects the benefits of risk diversification, but no clear effect on the selection of riskier projects can be individuated since diversifiable risks are variously distributed among projects and depend on the composition of the portfolio. Moreover, the fact that riskier investments offer higher returns is not obvious in the case of scientific projects (as discussed in sec. 3.3), while the meaning of 'HRHR' is ambiguous and depends on the assumptions made about the statistical properties of project returns. A quantitative and empirical effort to better assess the distribution of scientific returns is needed; understanding the origin of the covariance between projects is necessary to evaluate how risk diversification actually works for granting agencies.

The rest of the paper is organized as follows. The concept of 'risk' and 'return' and the mechanism of risk diversification for scientific research is investigated in sec. 2. Sec. 3 discusses the composition of the portfolio of a granting agency and the effects of risk diversification '*as if*' scientific research projects were typical financial assets and the selection mechanisms in granting agencies were the same as in a mutual fund (same mission, same preferences, same incentives). Sec. 4 deals with some (institutional and technical) specifics of science funding. Sec. 5 proposes several recommendations for improving, from a dynamic perspective, the coherence and effectiveness of agency procedures. Sec. 6 concludes.

2 Returns and risk diversification in scientific research

In this section I deal with two correlated issues: on one hand I want to introduce a 'micro-foundation' for risk management in granting agencies, i.e. the utility and the scope of risk diversification in scientific research project portfolios; on the other hand, it is worth formally defining what HRHR research really represents in order to understand how valuable risk management should actually promote breakthrough research while eliminating unproductive risk.
The scientific literature, with few exceptions, treats the distribution of returns in scientific research in a very qualitative way. Unfortunately, this is not sufficient when dealing with risk management in a portfolio. A working – albeit general – definition of the meaning of 'return' and 'risk' in scientific research is therefore required. In this section, I will firstly review the scant literature available before making some more formal proposals.

2.1 Literature on the risk and return of scientific research

Scholarly contributions to the literature on scientific research sometimes discuss the nature of its returns and underlying sources of uncertainty. In this respect, specific proxies for measuring returns (in particular when a relevant part of the output of the project is non-monetary) are proposed, where bibliometric indicators (papers published, patents filed, citations received, ...) are most privileged.[3] However, there has been relatively little systematic research and only some implicit and indirect evidence – empirical or conceptual – of the statistical distribution properties of the returns from scientific research.

Strictly referring to scientific research, the academic idea of risk is often associated with the variously defined concept of '*novelty*'. Azoulay et al. (2011) compare the research output of various sets of projects: they find that investigators using more novel keywords produce more hits and more flops, i.e. outcomes in the tails of the distribution. This is a perspective of risk that evokes *symmetric* deviations from an average outcome. Similarly, Wang et al. (2017) identify novelty as a crucial dimension of research risk, showing that novel papers are more likely to be top-cited papers but at the same time are also riskier, reflected by a higher variance in citation performance. Gans and Murray (2012, p. 85) distinguish between the short and long term returns of scientific research "*... projects differ in terms of their potential immediate social benefit, ... and their potential present value of future scientific benefits ...*" to investigate the complementarity/substitutability of private and public funding. They model short and long term returns as uniformly, independently, and identically distributed. Referring to the case of uncertain R&D projects in private firms, Huchzermeier and Loch (2001, p. 87) identify "*five types of operational uncertainty ... performance, cost, time, market requirement, market payoff* ", which usefully highlight how outcomes can be uncertain in many respects, not just in their disciplinary approach.[4]

A completely different approach is proposed by Vilkkumaa et al. (2014): they assume that the return

[3] "*The rare nature of transformational outcomes handicaps formal program evaluation, which is geared toward measuring short term outputs like patents and publications, or even medium-term outputs like commercialization activity* " (Azoulay et al., 2019, p. 70); "*... the Sloan Foundation specifically requests tangible outputs (such as number of students whose training or careers are affected, data collected, scientific papers produced) and outcomes (such as new knowledge, institutional strengthening, etc.) or other measures of success including big sales of a book, a prize awarded for research, a government grant to continue the project, web traffic, high enrolments, better salaries, etc. in evaluating grant effectiveness* " Gans and Murray (2012, p. 75).

[4] "*Research creates new knowledge or information, and so by definition, it is impossible to predict the result of a particular project. Furthermore, even if a project meets its technical goals, there is great uncertainty in how the research will be used and what the impact of that application will be*" (Azoulay et al., 2019, p. 80)

of scientific research is not stochastic, uncertainty only depends on noisy evaluations of the panellists.[5] Their choice is aimed at focusing on the benefits of repeated selective evaluation: "... *Bayesian modelling allows us to study how resources should be spent in order to derive more accurate value estimates, for instance by acquiring additional independent evaluations about some selected projects ... the expected portfolio value can be increased significantly by re-evaluating only a small fraction of projects (rather than spending resources on the re-evaluation of all projects, which can be very costly)*" (Vilkkumaa et al., 2014, p. 772). In this perspective 'risk management' only concerns the process of information acquisition; I return to this point in sec. 5.

2.2 Measuring the risk and return of scientific research

In this section I discuss a working approach to the definition of risk in scientific research. In this regard, I build on the work of Franzoni and Stephan (2021) who provide a very detailed analysis, first distinguishing between '*proper risk*' and '*Knightian uncertainty*' (or '*ambiguity*'), and then proposing a taxonomy for the components of risk in scientific projects (see their sec. 4).

Noisy evaluations of the projects (ambiguity) and heterogeneous evaluations among reviewers introduce specific and interesting sources of uncertainty that Franzoni and Stephan (2021) widely discuss. Instead, I focus here neither on the reliability of the opinions of the experts nor on the optimal aggregation of their evaluations into single scores.

In this paper, in general, a scientific project is a venture that can be launched thanks to the fixed investment of a given monetary sum. Except for what is discussed in sec. 5, investments (outflows) are considered as sunk costs concentrated at the beginning of the life of the project. The return of the project is a random variable with known (i.e. as mentioned, I do not deal with ambiguity) distribution. The assumption is then that returns can be measured; I do not discuss the nature of the unit of measurement.[6]

Uncertainty (risk) is determined by the fact that the research can alternately uncover many different *states of the world* (or *scenarios*), each with a different return (social value).

Assuming a discrete distribution of the J possible states of the world – where the generic state j is

[5] This condition directly implies the emergence of a winner's curse: "*if the ex ante value estimates are unbiased, those projects whose value has been overestimated most are more likely to be selected. It therefore follows that the realized ex post value of the selected portfolio often falls short of what is suggested by the ex ante estimates of the projects that it contains. In other words, the decision maker is likely to experience post-decision disappointment*" (Vilkkumaa et al., 2014, p. 772)

[6] In line with the financial approach assumed below, in the whole paper I consider a *normalized* measure of returns, i.e. the return/utility per investment unit.

realised with probability p_j and yields I_j –, the return profile of a single project is unambiguously represented by the prospect $Y = \{I_1, I_2, ...I_j, ...I_J\,; p_1, p_2, ...p_j, ...p_J\}$. The expected return of the project is $\mu = \Sigma_j\, p_j I_j$; the risk (volatility) of the project is reasonably captured by a measure of the dispersion around μ of the returns in the different scenarios, usually the standard deviation of the distribution. The expected utility of the project is defined as $EU\,(Y) = \Sigma_j\, p_j u(I_j)$ where $u' > 0$; the case of a concave $u\,(\cdot)$ simply displays risk-aversion, i.e. fixing the level of μ the expected utility increases as the volatility decreases.

As per Franzoni and Stephan (2021), a specific project can have many outcomes because even if every research project normally has a specific target, different secondary findings can be obtained as by-products of the main quest. The combination of successes/failures of primary/secondary goals can already determine a large number of possible scenarios. Moreover, each of these outcomes can be obtained taking more or less time (and thus, for example, being first innovators or not), with more limited/general results, affordable/expensive technical solutions, etc., thus representing still different scenarios. It is easy to conceive a situation of *infinite* states of the world, with returns I continuously distributed over the support $[\underline{I}, \bar{I}]$ with density $p(I)$. The prospect then becomes $Y = \{I \in [\underline{I}, \bar{I}]; p(I)\}$, and the expected return and utility modify accordingly.

As for the probability of each scenario, Franzoni and Stephan (2021, p. 14) suggest that this will depend on various components, properly aggregated: *methodological* uncertainty (further divided into epistemic, technical and organizational) and *natural* uncertainty. Other sources of uncertainty can be found in the features of the research group (uncertain competencies, creativity, health, synergy, ...) and the relevant scientific community (evolution of the state of the art, quality of competing research projects, ...).

In conclusion, very different return distributions can be individuated for different scientific research projects. However, a general discussion on the properties of a few paradigmatic cases is beneficial for i) understanding the mechanisms and the opportunities offered by rational and explicit portfolio management, and ii) properly defining what 'HRHR research' actually means; I argue that in common parlance the term 'HRHR research' is used ambiguously. Moreover, we are interested in understanding whether giving up the '*one-to-one approach*' should favour the financing of HRHR projects, in addition to the (small or large) benefits of diversification.

The extreme paradigmatic cases that I propose are two:

1. **Symmetric and continuous distribution** – The return variable is continuously and

symmetrically distributed around the expected return of the project. The risk of the venture is adequately captured by measuring the dispersion of the returns around the mean, the standard deviation (or *volatility*) σ.

Referring to the discussion above, *continuity* can be guaranteed by the combination of the multiple findings that the research can uncover, and the high sensitivity of the value of those findings to the underlying various conditions; *symmetry* is determined by a situation where the various conditions that determine uncertainty can move the expected value of the findings up and down with balanced magnitudes of the deviation (e.g., the research team could be tight-knit, organised and lively or the opposite with similar probability).

In this case, the idea of HRHR is reasonably associated with projects that present higher values both for μ and for σ, see fig. 1.[7] The statistical properties of portfolios in terms of their expected returns and volatilities as a function of expected return and volatility of the selected single projects follows the rules illustrated in sec. 3, and the preferences of decision makers are fully represented by a risk/return trade-off (the elasticity of substitution between the two in the utility function).

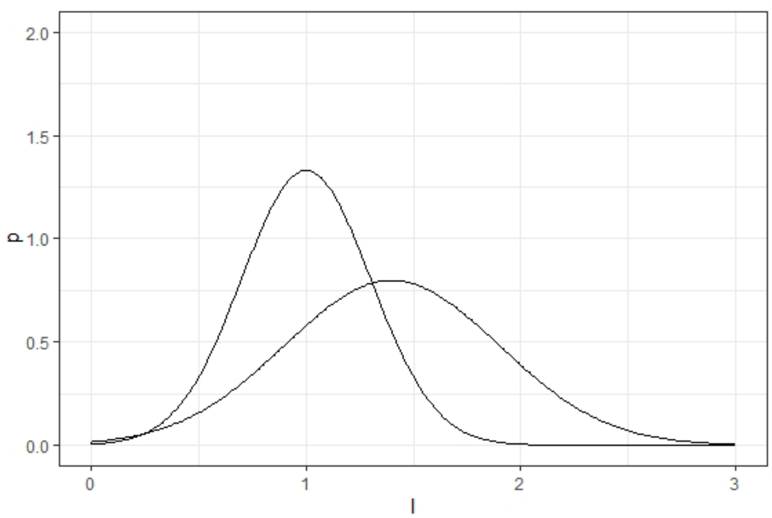

Figure 1: Two projects with symmetric distribution of returns: LRLR (left) and HRHR (right)

2. **Skewed and discrete distribution** – The generic project has a binomial return: success/failure. In case of success (with probability $p_S < 0.5$) the return is I_S; failure yields 0

[7] The implicit comparison of HRHR projects is with low-risk/low-return projects; in a financial framework, low-risk/high-return and high-risk/low-return projects are dominant/dominated, respectively, and are selected/discarded by risk-averse individuals immediately, with no need for portfolio considerations. For more in general, refer to the discussion in sec. 3.3.

with probability $1 - p_S$. In this case, the distribution can still be fully represented by the two parameters, p and I, or, equivalently, μ and σ, which are still the expected return of the project ($p_S I_S$) and the corresponding standard deviation ($I_S [p_S (1 - p_S)]^{½}$), but σ is not a good measure of risk anymore because it does not represent a symmetric magnitude of deviations from the expected outcome (the distribution is now *skewed*). When the expected value is very close to zero (i.e. the downside deviation from the mean is very likely and very small, while the upside deviation is very unlikely and very large), the main character of these projects (*high-skewness*, **HS**, projects) is better captured by their outstanding skewness more than by the dispersion of the returns around the expected value. In common parlance, the idea of HRHR projects is often attributed to HS projects.

Referring to the discussion above, a distribution of returns of this type can emerge when projects can end in a few possible scenarios/a narrower mission. For example, as Franzoni and Stephan (2021, p. 13) suggest, when no secondary findings are expected and/or when the research aims at empirically testing a hypothesis. In the latter case, there are often two outcomes: the hypothesis is successfully verified or not.

Fig. 2 represents two projects of this type (A and B). The two projects exemplify an interesting case with a 'conflicting' situation. Projects A and B are characterised by (limited) differences in terms of their expected return and volatility (p_A = 25%, $I_A = 1$, so that $\mu_A = 0.25$ and $\sigma_A = 0.43$, while p_B = 10%, $I_B = 2$, so that $\mu_B = 0.20$ and $\sigma_B = 0.60$); conversely, their skewness is very different: 1.15 for project A and 2.67 for project B (the skewness is 0 when the distribution is symmetric).

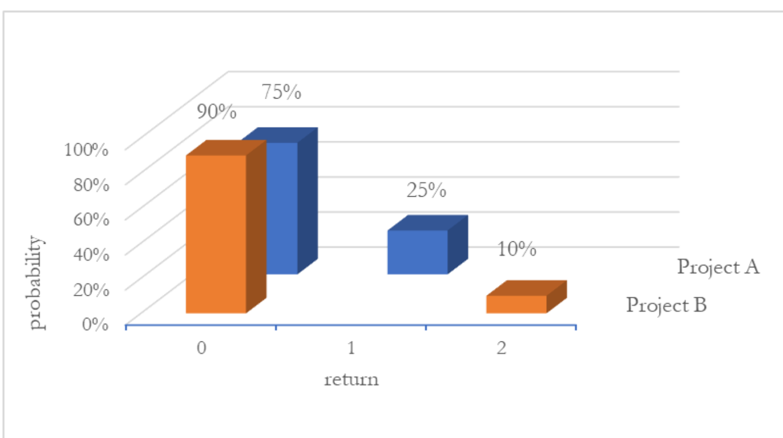

Figure 2: Two projects with skewed distribution of returns: high return-low volatility-low skewness (Project A) and low return-high volatility-high skewness (Project B)

When faced with a choice between *A* and *B,* a 'traditional' risk-averse individual (i.e. one that only evaluates the trade-off between expected return and volatility) consistently favours *A*, but a preference for positive skewness could suggest rationally investing in the HS project *B* instead of in *A*. Empirical evidence demonstrates that many decision makers present specific '*preferences for positive skewness*' which are not fully represented by a return/volatility trade-off.[8] The statistical properties of the portfolios composed of assets with skewed returns do not follow the rules I will illustrate in sec. 3.

To understand the role and the nature of risk management in a portfolio, starting from the two paradigmatic cases just presented, it is important to underline that 'portfolio effects' are actually determined by two different mechanisms:

- **proper risk diversification**, i.e. the possibility that some deviations from the expected value of single projects can neutralise each other when they have opposite signs. This naturally happens more effectively the more deviations are symmetrically distributed around the expected value and the less deviations are positively correlated (which happens when most of the risk is *systematic*);
- **size effect**, i.e. the effect of the law of large numbers, as the results obtained from a large number of trials become closer to the expected value.

In section 3, below, proper risk diversification – which prevails when the distribution of the returns of scientific projects is continuous and symmetric – is examined in depth. The case of discrete and skewed distributions is addressed in section 4.2.

3 Risky research projects and Markowitz

I start by simply embracing the traditional 'financial approach': scientific projects are treated as financial securities whose return is the expected financial performance of a continuous monetary

[8] Mitton and Vorkink (2007, p. 1255) claim that "*A substantial body of research documents that investors commonly hold portfolios made up of far fewer securities than are necessary to eliminate idiosyncratic risk ...*" and "*... investors may consciously choose to remain under-diversified to increase the likelihood of extreme positive returns, or in other words, to capture higher levels of skewness in their portfolios. Under fairly general conditions, (Arditti, 1967) and (Scott and Horvath, 1980) demonstrate that investors prefer positive skewness in return distributions. (Simkowitz and Beedles, 1978) and (Conine and Tamarkin, 1981) argue that when the third moment of the return distribution is taken into consideration, investors may optimally choose to remain underdiversified* ". If this is the case, one might argue that part of the unsatisfactory appreciation of HRHR by reviewers could be explained by a selection entrusted to risk-averse agents by skewness-prone principals.

variable (the market price of the security). Accordingly, the agency panel assumes the role and the goals that portfolio theory attributes to the mutual fund management team: the aggregation of single assets into an efficient (here 'efficient' means 'utility maximiser') portfolio.

In their "*Lack of portfolio approach*" section, Franzoni et al. (2021) claim that the 'one-by-one practice' usually adopted by granting agencies goes against the recommendations in financial literature, thus "*foregoing any advantages from using the portfolio to diversify away the risk* " (p. 16). As a consequence of this 'one by one' approach and the corresponding lack of interest in the composition of the entire portfolio, granting agencies ("... *to the extent that they are risk-averse*", p. 16) will fund fewer risky projects and then curtail the amount of HRHR research they finance. From a 'financial' perspective this latter claim is not true: several underlying assumptions must be investigated and the responsibilities of risk-aversion, on one side, and of the 'one by one' approach on the other, need to be correctly attributed.

As mentioned, in this section the scientific projects submitted for funding are treated as 'traditional' financial assets, i.e. their return is assumed to be a normally distributed random variable. Returns are stochastic but the parameters of the distributions are not uncertain. I will consider the possibility of correlated projects, as described by a covariance matrix. This perspective is in line with the 'case 1' proposed in sec. 2.2. In fact, returns here are assumed to be continuous and symmetric, with the additional assumption that they are normally distributed, as the evidence in the literature suggests for financial securities. Considering scientific projects as assets whose return is normally distributed, two parameters are sufficient to unambiguously characterise the distribution of the random variable: the expected performance and the risk are captured by the mean return μ and the volatility σ, respectively.

It is easy to translate this setting into the quantitative language of 'modern' portfolio theory in the vein of Markowitz (1952, 1959). Less obvious is how to assess the distance between the Markowitz framework and the case of granting agencies – both in terms of institutional and quantitative assumptions – which is the issue I deal with in sec. 4.

3.1 Risk diversification for research projects with normally distributed returns

Firstly I will analyse the case in which an agency panel has been appointed to select a portfolio of projects out of the set M of submitted proposals; the generic i project is characterised by expected return μ_i and volatility σ_i. The cost for funding each project is a constant amount, say one investment unit; every project can be financed entirely or not financed at all.[9] Assume first that the returns of the single projects are statistically independent; thanks to this (strong) assumption, if the agency selects the subset N composed of n projects ($N \subseteq M$), the expected return (μ_p) and the volatility (σ_p) of the portfolio are:

$$\mu_P = \frac{1}{n}\sum_{j\in N}\mu_j \tag{1}$$

$$\sigma_P = \frac{1}{n}\sqrt{\sum_{j\in N}\sigma_j^2} \tag{2}$$

An optimal portfolio selection (given the independence assumption) can then be achieved by estimating just two parameters (μ and σ) for every single project. The optimality for risk-averse decision makers depends on the corresponding level of risk tolerance (i.e. the ideal risk/return trade-off). Eq. 2 shows that – unlike the return – volatility is sub-additive, thus producing *risk diversification.*

In general, the higher n is, the more diversification is produced. However, the higher n is, the lower the average quality of the portfolio is because better projects (i.e. those with a better risk-return trade-off) are selected first. The optimal portfolio must therefore balance a good average quality of projects with a significant reduction of risk through diversification.[10]

It is interesting to understand how much of the risk can be neutralised simply by increasing the number of projects included in the portfolio, i.e. the magnitude of the diversification effect determined by the size of the portfolio. Of course, such a measure depends on the statistical properties of the distribution of the returns of each project. In the simple case above (statistical independence of the returns), if we also assume the n projects to be equal (same μ and σ), the risk of the portfolio $\sigma_p(n)$ is reduced by a factor $\Delta\sigma = 1/n^{1/2} - 1$ compared with the risk of the single project. This relationship is reported in fig. 3. For example, if $n = 2$ the risk is reduced by 29%, if $n = 5$ the risk is reduced by 55%, if $n = 10$ the risk is reduced by 68%.

[9] Of course, this assumption goes against the usual premise in portfolio theory that financial securities are perfectly divisible, and also against the widespread discretion of agency panels to scale grants.

[10] If the quality of applications is high, the budget constraint defines n, i.e. all available funds are used.

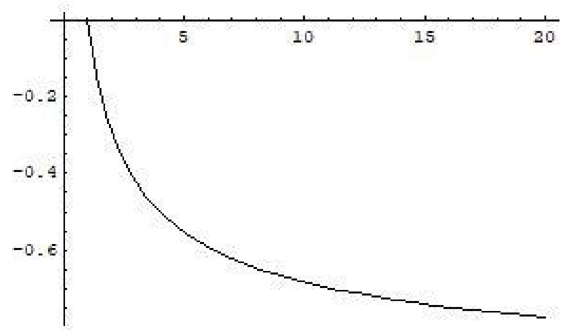

Figure 3: Diversification effect for *n* equal non-correlated assets with normal returns

In a more realistic case, however, project returns are correlated, even if imperfectly. For the sake of simplicity, let's assume that the set of available projects is composed of just three projects, *A*, *B*, and *C*, and the granting agency is subject to a budget constraint (one, two, or three investment units). Given the normality assumption, every project or portfolio is unequivocally individuated by a point in the risk-return space. The relevant parameters of their returns are summarised in tab. 1.

	Expected return (μ)	Standard deviation (σ)	Correlation matrix (ϱ)		
			A	*B*	*C*
A	1.00	0.30	1.0	0.9	0.3
B	0.95	0.25		1.0	0.0
C	0.70	0.22			1.0

Table 1: The moments of the returns of the three projects *A*, *B* and *C*

A, *B* and *C* are then the three red points in fig. 4, where the returns (μ) are reported on the vertical axis and the volatilities (σ) on the horizontal one. In a 'one-by-one' perspective, *A* is the project that promises the highest return but also the highest risk (can we define it as an example of an HRHR project?), while *C* seems to be the perfect example of an uninteresting project, much less 'profitable' and not that safe.

The (convex) iso-utility curves reported in the figure represent the risk-averse attitude of the decision maker. Utility obviously increases in a north-western direction. Steeper curves represent less risk-tolerant (more risk-averse) decision makers.

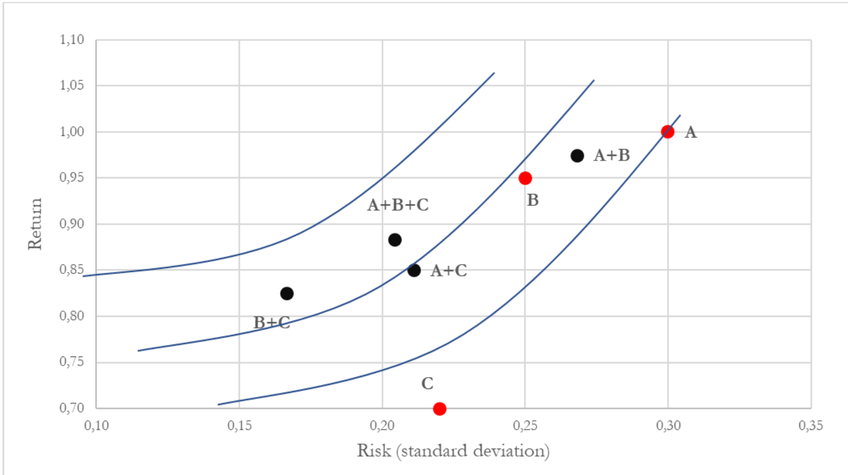

Figure 4: Projects and portfolios in the risk-return space

A 'one-by-one' approach would be forced if the budget of the agency were just one unit: in this case, the panellists must select *A* or *B* or *C*. The slope of iso-utility curves in the figure is such that *B* (mid risk, mid return) is selected. But a flatter (steeper) set of iso-utility curves (higher(lower) risk tolerance) would even have selected *A* or *C*. In this sense, it is the risk attitude that determines the risk/return trade-off that drives the decision: *any* project might be selected, provided that risk and returns are positively correlated.[11] If projects are sorted based only on mean numeric ratings received from peer reviewers (a measure of expected returns, the first column in tab. 1), *A* is selected: this means that risk is neglected, i.e. decision makers are risk-indifferent (the iso-utility curves are horizontal).[12] [13]

If the financial budget of the agency is two or three units, the properties of the returns of every portfolio must be calculated. In short, the return of a generic portfolio is still normally distributed, with a return that is the mean of the returns of the composing assets, while the volatility is still sub-additive.[14] Fig. 4 accordingly shows the position of every possible portfolio.

11 This is the case in our example: higher risk always corresponds to a higher return. This necessarily happens in a financial market at the equilibrium. Pleskac and Hertwig (2014) argue that this is true in general for 'ecological' reasons. This is not necessarily true in the case of scientific project financing. Refer to the discussion in sec. 3.3.
12 Contrary to common belief, in this setting a 'one-by-one' approach would *always* favour HRHR projects.
13 Linton (2016) implicitly suggests that risk preferences can somehow be taken into consideration by a suitable 'adjustment' of the scores. This is a questionable approach to risk management; in any case, by definition a 'one-to-one' approach, whatever the procedure, neglects covariances, which are the main ingredient of a correct portfolio approach.
14 Eq. 2 can be generalised in the presence of non-null correlations among projects: the volatility of a portfolio σ_p composed of n assets becomes:

$$\sigma_P = \frac{1}{n}\sqrt{\sum_{i=1}^{n}\sum_{j=1}^{n}\sigma_{ij}} \tag{3}$$

A 'one-by-one approach' (and risk-aversion) would rank projects along with a measure of risk/returns (in our figure, qualitatively, $B \succ A \succ C$), but proper attention to covariances among projects easily indicates $B + C$ ($\sim A + B + C$) $\succ A + C \succ A + B$. Of course, the case proposed depends on the discretionary value attributed to the parameters, but it is sufficient to show that neglecting covariances ('one-by-one approach') can lead to inefficient allocations of funds. Project *A* with the highest μ is selected only in a 'one-by-one' approach (and/or when investors are almost risk-indifferent).

The Markowitz (and followers) theory of investments simply states that by combining risky projects in a portfolio, part of the risk can be eliminated. The risk that can be eliminated, the diversifiable risk, *must* be eliminated because it is a sort of *unprofitable* risk since it does not contribute to the increase of expected returns. The optimal portfolio ($B + C$ in our example) is the one that better eliminates the diversifiable risk. From a one-by-one perspective, *C* is an inferior project (significantly lower return and comparable risk), but it adds significant value to the portfolio since it helps to diversify a large part of the whole risk thanks to its independence from the other two projects, which instead are fairly correlated. Also, notice that adding *A* to the $B + C$ portfolio does not increase the utility (compare in fig. 4 the utility associated to $B + C$ and $A + B + C$).

The following case is an example of the situation described so far. Three different projects share a single mission: *A* and *B* use a very similar methodology (maybe the most promising one). The opinion of the experts is unanimously high. A third proposed project *C* follows a different methodology for the same goal which has been classified by the experts as of low-medium merit. However, when considering the effects of the compositions of projects on the aggregate risk, things change. The risk that is diversified within the portfolio can be measured.[15] When combining *B* and *C*, half of the risk of both projects is diversified away; adding project *A* increases the return by 0.05 but also the risk by slightly less than 0.04. Given the risk-tolerance expressed by the iso-utility curves, the combined effect is of little or no incremental utility.

---

where $\sigma_{ij} = \sigma_i \sigma_j \varrho_{ij}$ is the covariance between *i* and *j*. When $i = j$, σ_{ij} is obviously the variance of the *i*-th asset (σ_i^2).

[15] The contribution of the particular asset *i* to the portfolio's risk is:

$$\sigma_{i,P} = \frac{1}{\sigma_P}\sqrt{\frac{\sum_{j=1}^{n} \sigma_{ij}}{n}} \qquad (4)$$

In other words, $\sigma_{i,P}$ (the non-diversifiable risk of *i*, obtained taking the first difference of eq. 3) is the part of the risk that *i* actually adds to the risk of the portfolio. The diversified risk is the difference between the whole risk of the asset and the non-diversifiable risk. As eq. 4 clearly shows, the share of diversifiable and non-diversifiable risk is determined by the composition of the portfolio.

3.2 Summing up

The examples outlined above can easily be extended to the general case with more projects (normally distributed returns and non-negative risk/return correlation) of different sizes. The main results are maintained, and in particular:

- the risk-indifferent decision maker selects projects in decreasing order of expected returns ($A - B - C$ in the example in fig. 4). He is only interested in the expected returns of each project (i.e. the mean numeric score received from the peer review panel) and he neglects the risk of the project (expressed by the volatility of the mean return). HRHR projects are selected first;

- 'one-by-one' selection in the presence of risk-aversion – provided that both the expected return and standard deviation of every single project are estimated – introduces a possibly different subjective rank (depending on the level of risk tolerance) based on the risk-return ratio ($B - A - C$ in the example in fig. 4); in any case, the composition of a portfolio based on this rank is not optimal since the opportunity to diversify part of the risk, which is a utility-increasing opportunity for risk-averse individuals, is ignored;

- the position of the different portfolios in the $\sigma - \mu$ space does not depend on the idiosyncratic attitude of the decision maker, while the choice of the optimal portfolio does depend on the level of risk-tolerance of the decision maker;

- if HRHR projects are those located in the north-eastern part of the $\sigma - \mu$ space, diversification does not guarantee that riskier projects are more likely to be selected; the composition of the optimal portfolio privileges those projects with a better trade-off between their return and their *marginal* contribution to the risk of the whole portfolio (i.e. after considering also their contribution to risk diversification); the benefits of risk diversification marginally decrease. For example, in the case in fig. 4, composing the portfolio through optimal risk management reduces the probability that the HRHR project will be selected;

- if HRHR projects are considered as poorly represented in the portfolio of the funding agency, this has to necessarily be attributed to a higher level of risk aversion of the decision maker; for example, an agency problem between more risk-averse panellists and less risk-averse agency.

The situation modelled by portfolio theory easily results in normative implications: correlation between projects must be taken into account, simple aggregated quality scores are not sufficient to obtain an optimal portfolio design, variances and covariances must be estimated and taken into account. This is simple to state, but difficult to arrange. The panel of experts will have to analyse the entire set of submitted proposals instead of focusing their attention on single ones, taking into consideration not only the expected outcomes, but also the sources of uncertainty and their combinations.

The more panels are focused on specific sub-disciplines that share risk factors and the more similar the proposed methodologies, the higher the portfolio risk will be, all else being equal. This situation raises new types of challenges for reviewers: as Franzoni et al. (2021, p. 25) put it: "... *Correlation between research paths, in and of itself, can be hard to determine, particularly when covering vastly different goals across different fields and with different research approaches*".

In conclusion, understanding the sources of the correlations between the returns of different projects become the main (empirical and conceptual) challenge for estimating the relevance of the benefits of the diversification techniques, and consequently proposing suitable adjustments to the prevailing selection procedures. Remember that if statistical independence between projects is postulated, then the one-to-one approach is not a mistaken procedure because it is the size of the portfolio that mainly drives the definition of aggregate risk. But correlation among projects is not an abstraction. We can conceive scientific, social and strategic reasons for justifying non-null covariances. On the scientific side, ventures using the same methodologies will present more correlated returns; for similar reasons, same scientific communities share the same knowledge base and enjoy the same findings and advancements; often they also share facilities and infrastructures (think of CERN, for example); on the contrary, different methodological approaches in the same portfolio decrease the aggregated risk more than proportionally.[16] From a social point of view, the

[16] Linton (2016) correctly suggests that alternative methodological approaches should be important from a reviewer perspective due to the possibility that alternative techniques and methods for the same mission might guarantee higher aggregated success probabilities (e.g. magnetic versus inertial in fusion research). The benefits obtained using alternative methods for solving specific scientific problems is well exemplified by the successes obtained by John Nash: "*Perhaps Nash's greatest mathematical work came from studying a mathematical puzzle that had been suggested to him by Louis Nirenberg. It concerned a major open problem concerning elliptic partial differential equations. Within a few months, Nash had solved the problem. It is thought that his work would have won him the Fields Medal – the most prestigious prize in maths, open only to those under 40 – had it not been solved at the same time by the Italian mathematician Ennio De Giorgi. The men used different methods, and were not aware of each other's work – the result is known as the Nash-De Giorgi theorem. One of the many amazing aspects of Nash's career was that he was not a specialist. Unlike almost all top mathematicians now, he worked on his own, and relished attacking famous open problems, often coming up with completely new ways of thinking*"(Bellos, 2015). In other words, funding both De Giorgi's and Nash's works would have optimally reduced the risk of not solving Nirenberg's puzzle.

scientific community is a network, and the proximity within the network creates higher probability of knowledge spillover (the corresponding literature is huge): the correlation in this sense could emerge from geographic and cultural proximity. Finally, a lot of team-specific sources of uncertainty (like those listed in the previous section) can be imagined as i.i.d. among different projects and thus determine lower values in the covariance matrix.

In conclusion, covariances in the setting of this section must be considered as crucial ingredients for correctly managing the risk, and neglecting them can determine severe mistakes in the allocation of funds. However, the relevance of risk diversification can significantly vary depending on the disciplines and the mission of the agency.

3.3 Risk/return correlation: financial assets vs. research projects

A final discussion in this section is devoted to questioning the assumption – which is taken for granted in a financial setting – that high-risk necessarily corresponds to a high-return. In general this is not true when considering research projects, even if with normal returns, due to their finite supply elasticity and the fact that here the risk is exogenous.

At the equilibrium, in an efficient financial market higher risks *must* be associated with higher returns. Imagine that in fig. 4 a fourth security is traded, but expected returns and risk are such that it is located in the south-eastern part of the graph (high-risk/low-return). No one would invest in such a security because other securities on the market – with higher returns and lower risk – are unquestionably preferred by risk-averse, or even risk-indifferent, investors. Remember that financial securities are supplied with unlimited capacity. In the absence of demand, the price of the *dominated* security decreases so that *the expected return increases and risk decreases*. In other words, the security moves north-west in the risk/return space until reaching a non-dominated position.

This case is in line with the traditional thinking that ventures/projects/investments with the highest potential also entail higher risk. However, the positive correlation between risk and return is confirmed only if risk and return are determined by human decisions (i.e. they are endogenous) and/or supply of dominant opportunity is not limited. Pleskac and Hertwig (2014) survey several alternative prospects and observe "... *a negative relationship between probabilities and outcomes, with larger rewards (payoffs) being associated with more risk (smaller probabilities). This negative relationship appears to take the form of a power function, especially in those environments in which probabilities and outcomes can be measured without much noise*" (Pleskac and Hertwig, 2014, p. 2004). They propose several cases and not surprisingly observe that the correlation is clear for roulette, life insurance and horse-race betting but much less so for the artificial insemination of dairy cows. In fact, they fail to observe that the rules of the game are determined by human hands under a participation constraint of risk-averse individuals for the former, but by nature for the latter.

The return on investment in real projects (as opposed to financial securities) is determined by the difference between revenues and operative costs (possibly normalised, for example, by costs). A firm weighing up various investment opportunities can optimally decide to invest in a portfolio of projects including 'dominated' investments[17] provided that the risk/return trade-off of this project lies above an appropriate threshold. Of course, non-dominated projects would be more appreciated, but outside the financial world they are not supplied with infinite elasticity.

Why should all this be relevant for our analysis? A scientific project is more like a business venture in this respect. The proposals submitted for funding will present their own risks and returns, with imperfect positive correlation. HRHR projects are obviously better than HRLR projects. Risk management – in the sense discussed below – can be very useful, but the relevant selecting rule simply requests the 'risk premium' to be sufficiently high, and this condition does not necessarily eliminate every 'lower return-higher risk project'.[18]

4 A granting agency is not a mutual fund

In this section I explore how far the analogy between science funding and financial investments holds, in particular when analysing incentives and mechanisms of risk diversification. Needless to say, a portfolio of scientific research projects differs in many respects from a portfolio of traded financial securities. While, to the best of our knowledge, the issue of diversification in portfolios of scientific research projects is almost new in economic literature, there is a long-standing literature that deals with the issue of risk diversification in VC portfolios (Norton and Tenenbaum, 1993; Cochrane, 2005; Ewens et al., 2013; Cressy et al., 2014) and in private R&D investments portfolios (Baysinger and Hoskisson, 1989; Garcia-Vega, 2006). At first sight, a scientific research project seems to have much more in common with R&D and VC than with traditional financial

[17] By *dominated* I mean with lower return and higher risk as compared to other projects in the portfolio.

[18] In other words, there is no inference of payoffs from probabilities or vice versa as suggested as a 'heuristic rule' by Pleskac and Hertwig (2014, p. 2001, 2002, 2006)"... *if payoffs and probabilities are interrelated, then this ecological property can permit the decision maker to infer hidden or unknown probability distributions from the payoffs themselves, thus easing the problem of making decisions under uncertainty ... people appear to use a simple heuristic, the risk–reward heuristic, to infer unknown probabilities from observable payoffs during decisions under uncertainty ... the risk–reward heuristic envisions that when faced with choice under uncertainty people infer that the probability of an event is negatively related with the magnitude of the payoffs*".

investments.[19] In common with VC or R&D, scientific projects are 'lumpy' i.e. they are investments that are either wholly selected or rejected. This differs from financial portfolio choices, where essentially any fractional amount of resources can be invested into any asset.

4.1 The nature of granting agencies: goals and attitude

4.1.1 Non-private returns and public goods: is risk diversification an issue for granting agencies?

Before addressing the validity of the technical assumptions of the portfolio theory in our context, I focus on the *institutional* specificity of 'portfolio asset allocation' for granting agencies. In this respect, the cornerstones of projects selection for granting agencies are the following:

- first, research projects in the portfolio of a funding agency produce a public (and not private, as in the other cases) value; this obvious observation conceals many subtle implications concerning how performances and risk diversification have to be evaluated;
- second, the structure of preferences for the returns of a scientific research portfolio cannot be mapped in monetary units. An agency granting funds in support of energy research projects, for example, might be not interested in 'optimal' risk diversification if this requires an investment in, say, medical research;[20]
- third, unlike the other types of portfolios I have mentioned, scientific research projects are never 'priced' on the market. Consequently, the ex-ante evaluation of projects must be entirely entrusted to experts, and the value produced cannot be precisely measured ex-post: this makes it more difficult to use incentivising devices to align the interests of agents (experts? agency administrators?) with those of the ultimate stakeholders (the sponsors? the collectivity?). The fact that these assets are not traded on the market also determines that

[19] Actually, in some respects the similarities are mixed: for example, unlike VC or R&D portfolios, but similar to a mutual fund, funding agencies are usually not required/allowed to actively manage their assets, i.e. as will be discussed in sec. 5, funded scientific projects are infrequently re-negotiated or modified during their life by agencies, while VC or R&D managers continuously adjust their ventures.

[20] Notice that specialisation in VC funds, in R&D investments, or even in mutual funds is mainly driven by the willingness to focus competencies, not by non-monetary preferences. "*While the finance literature argues that risk reduction is a primary benefit of diversification (Markowitz, 1991), its impact on portfolio building in the context of active managers is particularly important in VC, given the significant investment risks involved in these types of investments*" (Buchner et al., 2017, p.520). In this sense, the benefits of diversification (risk reduction) are traded-off against the corresponding costs (lack of specialisation).

risks and returns are not endogenous (see the discussion in sec. 3.3).

"*Each funded proposal must be of importance to society (House of Representatives, 2009) ... the desire for many policy makers* [is] *to have a clear path between research and the creation of economic and societal value*" (Linton, 2016, p. 1936). After considering the institutional setting sketched above, one may wonder whether risk diversification is an issue in our setting. Let's assume funding agencies are risk-averse (I will return to this point in sec. 4.1.2) and that they are able to identify and measure the return, risk and covariances of every single project. Let's take the extreme perspective of funding agencies solely aimed at contributing to the provision of useful scientific knowledge on behalf of an economic system, a country, human society, ... and assume the perspective of the previous example in fig. 4. What is socially optimal?

If the available amount of funds is, for example, two investment units, B and C must be financed.[21] Does it make any difference whether the two projects are both funded by a single granting agency or by two different agencies? Given that the projects contribute to the public good, the question of who is providing the funds seems to be irrelevant: if *B* and *C* are financed, the diversifiable risk gets diversified at a systemic level, independently of the financing structure.

Notice that such 'financing structure irrelevance' does not exclude the relevance of diversification matters at a single agency level. If funding agencies (or the collectivity) are risk-averse, they will still want as much risk as possible to be eliminated through diversification *from the portfolio of all the portfolios* (i.e. of all the different funding agencies) contributing to the provision of the public good. Consequently, in principle any agency should rank projects considering their *risk-adjusted returns*, like in eq. 4. This perspective is not different from the perspective of a financial intermediary who knows that the portfolio he is managing will be combined by the final investor, so that he is not necessarily asked to diversify (all) the risks, but when evaluating the assets he manages he must consider that someone else will further diversify the remaining diversifiable risk.[22]

---

[21] Another relevant difference that I won't stress in this work is that in the traditional theory of financial markets the supply of both investment opportunities and capital financing is perfectly elastic, so that every available asset ends up in the 'market portfolio' at an appropriate price. Here, scientific projects submitted and funding resources are both limited. However, project rationing is reasonably expected (funds can be considered scarcer than projects). In any case, even if projects are correctly ranked and selected, social optimality is not guaranteed: in an efficient financial market, returns are endogenous and determined at the equilibrium by the interaction of a crowd of price-taker operators; in the world of science funding one could expect that proposals are financed until funds are exhausted so that the 'productivity' of the funds invested is determined by the productivity of the marginal project. Is that level of productivity the optimal one? More generally, how is the socially optimal threshold of productivity determined conceptually?

[22] As an example, the manager of a mutual fund specialising in Italian stock denominated in euros is usually not asked to diversify currency and country risks: these risks will be removed by the final investor combining this fund with other funds in dollars, yen, etc. In any case, he must understand that the market will not remunerate currency risks as they eventually can – or rather, must – be eliminated.

The issues discussed in this section, however, have limited practical application. The idea that agencies' decision makers must care about risk diversification, ranking their projects through a measure of risk-adjusted expected returns, cannot be easily implemented. The difference between diversifiable and systematic risk in scientific projects is conceptually complicated.

A final concern has been raised by Franzoni et al. (2021, p. 25) when suggesting that "... *there is the question of fairness: in building a portfolio approach some proposals may have to be eliminated in an effort to balance or de-risk the portfolio*". Their point is that taking into account risk diversification might determine that 'better' projects (in terms of higher scores received from the peer review panel) could be discarded in favour of 'worse' ones, just because the 'worse' projects are less correlated with the rest of the portfolio than the 'better' ones. Correct diversification, from a technical point of view, could indeed determine rankings of projects which are not based on our traditional notion of merit, and this could be a deontological problem.[23] On the other hand, defining a 'new' concept of merit could introduce more efficient incentive schemes for scientists.

4.1.2 Why do granting agencies behave as if they were risk-averse agents?

The easy explanation for risk-aversion simply derives from the assumption that decision makers maximise the expected utility of a concave utility-of-wealth function. Whether the utility-of-wealth function for a funding agency is concave or not depends on the (social) marginal utility of the funds employed in scientific research, which is reasonably decreasing.

Moreover, provided that the procedures they follow inefficiently diversify the risk, risk-averse decision makers in funding agencies care more about risk than they should. In other words, the risk borne (or at least perceived) at portfolio level is higher than it should be, and this sees the agency act as if it is more risk-averse than it actually is. The supposed bias of funding agencies in favour of less risky projects could then be explained by this effect, even abstracting from traditional explanations (Franzoni et al., 2021).

If we relax the assumption that the utility of a funding agency is just a function of the amount of public good it provides, our evaluation slightly changes. Of course, single agencies must comply with accountability constraints; moreover, the endowment of funds

[23] For example, the peer review system used by NIH – as legally required through sections 406 and 492 of the PHS Act – is aimed at providing "... *a fair and objective review process in the overall interest of science*".

– in a dynamic setting – reasonably depends on past results. In this perspective, considering a trade-off between returns and 'locally' adjusted risk could make sense.[24]

Finally, decisions are taken by (risk-averse) individuals, whose utility is only imperfectly aligned with the utility of the institutions they are serving. In the literature, moral hazard (i.e. misalignment of objectives) is taken as ubiquitous among people involved in portfolio management activities.[25] Franzoni et al. (2021, sec. 3 and 4), and the literature cited there, provide theoretical and empirical arguments on the idea that the risk-aversion of the individuals taking decisions within granting organisations (either outsiders or insiders) determines the bias against risky research in the competition for funding. The unanimous opinion in this literature that misalignment of incentives determines the insufficient allocation of funds in 'more risky research' makes sense. However, precisely because of this, risk management, and in particular risk diversification, should be considered as an attractive practice for eliminating unproductive and/or misallocated risk.

In summary, risk in scientific research is a concept that still suffers from a considerable degree of vagueness. The analogies arising from the financial literature should be viewed with caution. Moreover, as discussed in sec. 2.2, the attitude of decision makers facing uncertainty is not completely described by their attitude towards volatility, because, unlike financial risk, scientific research risk is often non-symmetric.

4.2 Risk diversification for research projects with skewed returns

If we admit the possibility of skewed distributions of project returns,[26] the analysis conducted in sec. 3 is not appropriate because just two moments (mean and variance) are not enough for defining

[24] By 'locally' adjusted risk I mean the risk that remains after diversification *within the portfolio* of the agency, not the risk that remains after diversification at systemic level.

[25] The way the interests of managers can be aligned with the attitude of the institution is a central issue. As mentioned, the utility of professional (mutual or VC) fund managers can be more easily aligned; the compensation of panellists cannot be based on performance. Furthermore, the special case of granting agencies generates conditions of particular interest: a large number of decisions are often assigned to external experts, who can hardly perceive "*risk culture, as the organization's propensity to take risks ... it is the perception that creates the culture, even more than any tangible and documented set of decisions or actions taken by organizational actors, because it is the perceptions that provide the cues to acceptable behavior*"(Bozeman and Kingsley, 1998, p. 111).

[26] If we presume that bibliometric indicators could serve as proxies for the returns of scientific research, support for this hypothesis is provided by the literature, which indicates that nearly all distributions in bibliometrics are skewed. For example, Seglen (1992) reviews rich evidence – since Lotka's pioneering work (Lotka, 1926) – that skewness is characteristic for the productivity distribution of scientists. However, the skewness of scientists' productivity distribution determines skewness in the distribution of the outcomes of scientific activities only under specific assumptions.

the distribution, and the concept of 'risk' must be redefined. Franzoni and Stephan (2021), for example, point out that the outcomes of research vary predominantly in the spectrum of gains and potentially lead to exceptional results, but also to no results. In other words, they stress that scientific returns are lower bounded, i.e. not symmetrically distributed. Examining VC portfolios, Buchner et al. (2017, p. 525) maintain that "... *an important drawback of the standard return volatility is that it treats positive and negative deviations from the mean return as equally undesirable risk. In contrast, downside volatility is a risk measure that accounts for asymmetric return distributions by considering only negative deviations from a pre-specified target return*". This point, which can be easily translated for research projects, simply indicates that in asymmetric distributions the risk is not a univariate variable.
Downside volatility is therefore a more plausible measure of the risk when returns are highly skewed, while upside volatility is a measure of the upside potential: "... *upside risk is generated by high volatility of returns on the upside (i.e., through very high returns); thus, upside risk is deemed 'good' risk. In contrast, downside risk is 'bad' risk, because it captures the volatility of losses (negative returns)*" Buchner et al. (2017, p. 525).

I have discussed in sec. 2.2 the micro-foundations for the hypothesis of a skewed distribution of scientific returns (our 'case 2'). The literature also offers several (indirect) elements for seriously taking into consideration the cases both of *positive* and *negative skew* for the distribution of returns of scientific projects. The latter hypothesis is less obvious and must be associated with specific selection procedures. In this sense, I deal in this section with positive skewness (i.e. the case of the examples in fig. 2). A brief outline of the negative skewness case will be proposed in sec. 5.

A positive skewness of the distribution of returns indicates that the mean is larger than the median. Buchner et al. (2017) claim that positive skewness is a common feature of VC returns. More generally, in their paper "*Technology policy for a world of skew-distributed outcomes*" (2000), Scherer and Harhoff provide rich empirical evidence from very different data sets of the pervasive emergence of positive skewness in the returns from variously declined 'technological ventures': patents (US, Harvard, German, universities), start-ups, IPOs of technology-based firms, and innovative products (drugs). They find evidence of positive skewness showing that the fraction of total revenues contributed by the top 10% of each sample set is always higher than 48%, often more than 80%. This striking evidence merits two observations.
First, in the analysis of Scherer and Harhoff (2000), returns are always monetary variables (royalties, sales, capital gains, etc.). This means that the returns are underestimated, in particular for what

concerns long-term returns and social returns. The point is of obvious interest; for what concerns our discussion, one may wonder whether the inclusion of social and long-term returns could make returns distribution more symmetric. The answer is dubious.[27] Fragmentary evidence (e.g. Mansfield et al., 1977) suggests that the social returns from private investments are similarly skew-distributed.

Second, besides empirical suggestions, the type of ventures surveyed by Scherer and Harhoff (2000) – and those we consider in this work – naturally support the idea of positive skewness. A normal (or at least symmetric) distribution has to do with an expected outcome (say, revenues from sales of uncertain amounts of goods or services, the production of uncertain amounts of specific goods – 'knowledge' in our case) and the possibility of random deviations (with zero-mean) from expectations. On the contrary, technological ventures – and, in general, scientific research – frequently recall stories of 'success vs. failure', with a much higher probability of disappointment. Taking it to the extreme, for example, empirical research aimed at testing a formal hypothesis can present pure binomially distributed returns.

In the positive skewness framework outlined in sec. 2.2, the effectiveness of *proper* risk diversification is much more limited because downside and upside volatility are very different in their magnitude and frequency. In particular,

- the risk of the portfolio cannot be composed as per traditional portfolio theory and placing projects and portfolios in the σ-μ space makes little sense for representing preferences and decisions. The matter is that relatively close points in the σ μ space can be associated with very different skewness of the distribution of the return of the underlying asset, and a preference (or an aversion) for skewness can't be represented;
- in this case, risk management must be mainly concerned with size effects ('*spray & pray*'); if success is really unlikely and determined by unpredictable factors, the assessment of the corresponding probability is greatly unreliable, and the correlation between different projects even more so.[28] If this is the case, a 'one by one' selection procedure can be considered a good second best. More generally, this is true also for fairly uncorrelated sets of scientific

[27] As for the distribution of long-term returns of bibliometric indicators, see Wang et al. (2013).

[28] The literature identifies numerous cognitive biases that influence experts' judgment of the importance of extreme risks, in particular for extreme bad (catastrophic) risks; in this regard see, for example, Yudkowsky (2008).

projects;[29]

- in this framework, a volatility-averse/skewness-indifferent decision maker will always discard HS projects. If the granting agency is skewness-prone this situation entails a new different type of agency problem;

- a dynamic perspective (as opposed to the traditional static one) should better disclose the beneficial contribution of diversification. In this setting, risk management is expected to significantly contribute to rebalancing the presence of incremental and breakthrough research in the choices of granting agencies.

4.3 HRHR projects when the distribution of returns is skewed

The effects of diversification when investments concern assets with positively skewed returns are not easy to grasp. With skewed returns only a few of the projects supported will pay off on a large scale: a first idea is that increasing the number of projects supported can also increase the probability that large returns from the relatively few successes will also cover the cost of the many less successful projects. This is the 'size effect' that I mentioned in the previous subsection. Of course, this effect is diluted the higher the covariance is among the returns of the different projects.

The idea of HRHR that emerges from the discussion above identifies breakthrough research with a situation of high variance of returns together with high positive skewness of the distribution: a very high probability of failure, but a very high reward in case of success. Postulating that HRHR projects are not only characterised by high variance but also by high positive skewness is not just pedantry, as the presence of skewness has significant effects on risk diversification practices: if outcome distributions are sufficiently skewed, even with large numbers of projects, it is very difficult to 'properly' *diversify* away risk through portfolio strategies. In this perspective, the evidence of under-investment in HRHR projects could be determined by *positive-skewness-aversion* more than by the classical *volatility-aversion*.

Before addressing the specificities of diversification when investments show skewed returns, it is important to recall one crucial point that associates scientific research to VC investments and R&D

[29] This is obviously true even in the case of a symmetric and continuous distribution, as discussed at the beginning of sec. 3.1: if projects are all statistically independent the benefits of portfolio composition are due to the simple size of the portfolio; the point is that the assumption of statistical independence seems more reasonable with discrete and/or skewed distributions: with a small number of states of the world, the probability that specific determinants simultaneously hit various projects is much less likely.

projects. Unlike what happens on efficient financial markets – where the supply of investment opportunities is perfectly elastic – granting agencies (and assimilated ventures) face a limited number of 'good' projects. In this sense, all else being equal, a strategy of diversification based on increasing the number of funded projects reduces the average quality of the portfolio. This situation has to be kept in the background of the discussion that follows as an element that is always capable of reducing the possible benefits of risk diversification.

Focusing on risk management strategies based on increasing the size of the portfolio, Scherer and Harhoff (2000) analyse the properties of alternative skewed distributions and simulate the behaviour of portfolios returns. They show that no clear conclusion can be drawn. With a Pareto–Levy distribution it is difficult or impossible – by increasing the size of portfolios – to achieve stable mean expectations and hence hedge against risk.[30] On the contrary, log-normal distributions exhibit better-behaved large-sample properties, i.e. they better support portfolio strategies. Notice that all of the experiments by Scherer and Harhoff are conducted by fixing an empirical (not disclosed) correlation of the returns of the sample. Scherer and Harhoff (2000, p. 563) conclude that in the presence of skewed returns *"attempts to achieve stable mean returns through feasible portfolio strategies are likely to yield at best middling success"*.

Buchner et al. (2017) empirically analyse a large sample of VC portfolios in order to investigate the relationship between portfolio diversification and performance. They claim that much of the literature has examined such a relationship (i.e. the costs and benefits of specialisation) without considering the simultaneous impact on funds' risks. Because VC funds exhibit highly skewed returns, they separately consider the effects of diversification on both upside and downside risk. They show that diversification among industries increases returns, thus arguing that "*If a VC manager diversifies more across industries, the fund will carry less industry-specific risk, thus enabling the manager to engage in riskier portfolio companies (i.e., portfolio companies with a high downside risk but also a high upside potential in terms of returns)*" (Buchner et al., 2017, p. 522). However, they also find a positive and significant relationship between industry diversification and downside volatility (i.e. higher industry diversification increases the likelihood of picking losers). This result can be explained by the specific nature of VC investments: higher diversification among industries implies minor specialisation. Consistent with this, Buchner et al. (2017) identify a different effect when diversification is across stages (not industries): in this case, higher

[30] "*The Pareto–Levy distribution has the unusual property that when* $a < 1$*, the weak law of large numbers fails to hold, so that neither the distribution's mean nor its variance is asymptotically finite. What this means in practical terms is that as one draws ever larger samples, there is an increasing probability that some extremely large observation will materialize, causing both the mean and the variance to explode upward rather than converging toward stable values*" (Scherer and Harhoff, 2000, p. 563).

diversification reduces downside volatility.

The works illustrated above are not devoted to the analysis of scientific research projects, but many hints can easily be transferred to our case of interest. In particular, industry- and stage-diversification can have immediate meaning also for granting agencies and the issue of specialisation sounds relevant as well.

5 Dynamic risk management in scientific research

The flourishing literature on scientific research funding cited in the previous sections describes funding programmes as variously directed in the expression of their objectives and the design of their selection procedures. However, most of the programmes – both public and private – are similarly and uniformly organised for what concerns the timing of investments, the length of the funding period and the relatively loose control during the lifetime of the project. The uniformity of these rules is puzzling and echoes a correspondent doubt raised by Lerner and Nanda (2020) referred to recurrent practices in the VC industry.[31] This situation is particularly critical for scientific research projects that face (as in the VC industry) huge uncertainty about their ultimate potential, even when ex-ante selection is conducted effectively and properly. The skewed nature of the returns discussed in sec. 2.2 and sec. 4.2 also indicates that most of the investments are doomed to fail, and the few extremely successful investments are hard to be predicted upfront. The information that is released during the active life of the projects can quickly reduce such uncertainty. In this sense, while emerging information (and staged financing, i.e. increasingly larger amounts of funding are allocated, depending on whether interim milestones are met) is widely used by venture capital investors (both as a method for refining information and as a governance/incentive tool), it is largely neglected by granting agencies where upfront financing is the prevailing practice.

A second related remark concerns the fact that portfolios of risky R&D projects, VC ventures and scientific research projects present '*option characteristics*', i.e. a space for 'managerial' flexibility allowing investors to adjust decisions under uncertainty. In particular, the opportunity within the horizon of planned activity to prematurely end a project that proves to be inadequate or to expand

[31] "*Since the early days of the industry, venture capital funds have been eight to ten years in length, with provisions for one or more one- to two-year extensions. Venture capitalists typically have five years in which to invest the capital and then are expected to use the remaining period to harvest their investments. Funds differ tremendously in their investment foci: from quick-hit social media businesses to long-gestating biotechnology projects ... there is tremendous variation in the maturation of firms in different industries. Certainly, within corporate research laboratories, great diversity across industries exists in terms of the typical project length. What explains the constancy of the venture fund lives?* " (Lerner and Nanda, 2020, p. 253).

a promising one. The opportunity to adjust the grant during the life of the project, and not in a lumpy way, could have obvious effects in terms of risk management. In the lexicon of the previous section the benefit of this approach can help eliminate a large share of downward risk while retaining upward risk. In other words, managing risk from a dynamic perspective can modify the profile of returns distribution by reducing the positive skewness or even obtaining *negative* skewness, as option portfolios normally present.

Van Bekkum et al. (2009) develop these ideas in the case of R&D portfolios.[32] They distinguish between unconditional and conditional (i.e. with staging) projects, and show that when projects are positively correlated, the overall portfolio risk for conditional projects is lower than for unconditional projects. "*Diversification is therefore less effective than one would initially expect from unconditional investments, and more weight should be placed on non-diversification arguments to motivate a portfolio of such projects, such as synergies and spillovers*" (Van Bekkum et al., 2009, p. 1151); "*in a portfolio of options, the option to abandon limits downside risk of the individual project, but complicates diversification and does not limit risk when portfolio correlation is negative. In line with Jensen's Inequality, we call this the 'convexity effect', which affects the diversification effect*" (Van Bekkum et al., 2009, p. 1152).

The 'conditional perspective' of portfolio risk management consequently moves the focus from 'diversification' to 'de-risking', commonly used for funding entrepreneurial projects in the VC industry, where the expression 'spray & pray' has been coined ((Lerner and Nanda, 2020)). Ewens et al. (2018) also reference the increase of spray & pray investment strategies in the VC industry, where financiers provide a small amount of funding and limited governance to a larger number of start-ups.[33][34]

Of course, with the caution discussed earlier, the positive effects of the diversification of non-systematic risk still contribute to optimally allocating funds.

What has been discussed so far can be inscribed in the usual taxonomy of types of real options (delay, abandon, contract, expand, switch). "*Midcourse actions during R&D projects to improve the*

---

[32] Huchzermeier and Loch (2001, p. 85) also maintain that "... *most investment decisions (and R&D projects in particular) are characterized by irreversibility and uncertainty about their future rewards: once money is spent, it cannot be recovered if the payoffs hoped for do not materialize. However, a firm usually has some leeway in the timing of the investment*".

[33] "... *retaining control rights over the research after allocating funds ... 'go/no-go' reviews ... projects rated well are likely to see their budgets increased*" (Azoulay et al., 2019, p. 82-83).

[34] The interesting case of financial support for vaccine projects discussed by Veugelers and Zachmann (2020) also indicates that the virtuous properties of stage financing can be further enhanced when later stages are also the most expensive (managing large clinical trials and regulatory approval procedures), thus allowing investors to 'buy' cheap and increasingly reliable information during the development of the projects on success probabilities.

performance of the product (or to correct its targeting to market needs) are commonly used. The availability of such improvement actions represents an additional source of option value" (Huchzermeier and Loch, 2001, p. 86).

The rich corresponding literature provides interesting tools for measuring the financial value of flexibility and assessing the opportunity of investing through proper capital budgeting techniques. This part of the story is not really interesting for our purposes. However, the setting proposed for interpreting conditional projects in a real option perspective can be conceptually useful for understanding that 'risk' (or, better, upside risk) is valuable if downside risk can be reduced or eliminated. "R*eal options theory has shown that uncertainty in the market payoff enhances the project value if management has the flexibility to respond to contingencies ... the question is whether this insight holds as well for uncertainty in the other value drivers ... increased variability in market payoffs, as well as in budgets, may indeed enhance the option value of managerial flexibility, consistent with option pricing theory*" (Huchzermeier and Loch, 2001, p. 87).

6 Concluding remarks

The prevailing procedure for selecting projects to be funded out of a set of proposals by simply ranking them by individual expected quality/return introduces two specific types of allocative inefficiencies when decision makers are risk-averse. On one side, neglecting the risk, funds can be granted to projects with lower risk-adjusted return. On the other side – even if risk is considered, but correlations among projects are neglected – funds are allocated without considering that in the final selected portfolios individual risks will be cut with variable proportions (that depend on the whole correlation matrix).

Avoiding these allocative inefficiencies requires a significant incremental evaluation effort.

The most restrictive set of reasonable assumptions requests the returns of the projects to be all statistically independent and symmetrically distributed around their expected value. In this case, an efficient selection can be obtained from the estimate of just two parameters per project: its expected quality and the corresponding risk, as expressed by the dispersion of possible scenarios around the average outcome.[35]

However, in more general cases, provided that projects are imperfectly correlated, the optimal composition of a portfolio requires the joint analysis of all submitted projects.

---

[35] A very simplified index for ranking projects in this case could be μ-$\beta\sigma$, where β captures a measure of risk tolerance, i.e. more risk-averse agents show higher β values.

Scientific research returns are often characterised by the presence of very long right tails, i.e. a positively-skewed distribution; in this framework traditional risk diversification is much less effective. Moreover, when dealing with scientific research projects, diversification through the increase in size of the portfolio – unlike in the case of financial investments – might determine an average quality decrease, due to the limited supply of high-quality projects.

One prospective solution to avoid the ineffectiveness of 'classical' risk management techniques borrowed from financial literature is to organise the provision of funds in sequential stages – thus creating an 'option value' – with later funding granted contingent on the information released by preliminary results. This 'conditional perspective' of portfolio risk management moves the focus from 'diversification' to 'de-risking'.

One final point we have covered in this paper is associated with the widespread concern about the bias of current selecting procedures against HRHR research. The idea that risk diversification (in the traditional understanding of the procedure) should increase the probability of funding HRHR projects is not supported by meaningful arguments. The opposite might also be true. From a traditional (financial) perspective diversification simply helps eliminate unprofitable risk. But this risk is distributed along with covariances, not variances. However, the meaning of HRHR research in common jargon seems to be associated with high volatility, but also with high positive skewness. The bias of granting agencies against HRHR research must then be interpreted in terms of skewness-aversion, not in terms of risk-aversion.

References

Arditti, F. D. (1967). Risk and the required return on equity. *The Journal of Finance*, 22(1):19–36.

Arrow, K. J. (1962). Economic welfare and the allocation of resources for invention. In National Bureau Committee for Economic Research, editor, *The rate and direction of inventive activity: economic and social factors*, pages 609–626. Princeton University Press, Princeton, NJ.

Azoulay, P., Fuchs, E. R., Goldstein, A. P., and Kearney, M. (2019). Funding breakthrough research: promises and challenges of the 'ARPA Model'. *Innovation Policy and the Economy*, 19(1):69–96.

Azoulay, P., Graff Zivin, J. S., and Manso, G. (2011). Incentives and creativity. Evidence from the academic life sciences. *The RAND Journal of Economics*, 42(3):527–554.

Baysinger, B. and Hoskisson, R. E. (1989). Diversification strategy and R&D intensity in multiproduct firms. *Academy of Management Journal*, 32(2):310–332.

Bellos, A. (2015). John Nash's unique approach produced huge leaps in economics and maths. *The*

Guardian, May 24th.

Bozeman, B. and Kingsley, G. K. (1998). Risk culture in public and private organizations. *Public Administration Review*, 58(2):109–118.

Buchner, A., Mohamed, A., and Schwienbacher, A. (2017). Diversification, risk, and returns in venture capital. *Journal of Business Venturing*, 32(5):519–535.

Cochrane, J. H. (2005). The risk and return of venture capital. *Journal of Financial Economics*, 75(1):3–52.

Conine, T. E. j. and Tamarkin, M. J. (1981). On diversification given asymmetry in returns. *The Journal of Finance*, 36(5):1143–1155.

Cressy, R., Malipiero, A., and Munari, F. (2014). Does VC fund diversification pay off? An empirical investigation of the effects of VC portfolio diversification on fund performance. *International Entrepreneurship and Management Journal*, 10(1):139–163.

Ewens, M., Jones, C. M., and Rhodes-Kropf, M. (2013). The price of diversifiable risk in venture capital and private equity. *The Review of Financial Studies*, 26(8):1854–1889.

Ewens, M., Nanda, R., and Rhodes-Kropf, M. (2018). Cost of experimentation and theevolution of venture capital. *Journal of Financial Economics*, 128(3):422–442.

Franzoni, C. and Stephan, P. (2021). Uncertainty and risk-taking in science: meaning, measurement and management. Working paper.

Franzoni, C., Stephan, P., and Veugelers, R. (2021). Funding risky research. Technical Report No. 28905, National Bureau of Economic Research.

Fuchs, E. R. (2010). Rethinking the role of the State in technology development: DARPA and the case for embedded network governance. *Research Policy*, 39(9):1133–1147.

Gans, J. S. and Murray, F. (2012). Funding scientific knowledge: selection, disclosure and the public-private portfolio. In Lerner, J. and Stern, S., editors, *The rate and directionof inventive activity revisited*, pages 51–103. University of Chicago Press, Chicago, IL.

Garcia-Vega, M. (2006). Does technological diversification promote innovation? an empirical analysis for European firms. *Research Policy*, 35(2):230–246.

House of Representatives (2009). Investing in high-risk, high reward research. Hearing before the subcommittee on Research and Science Education Serial No.111-55, The Congress, Committee on Science and Technology, Washington DC.

Huchzermeier, A. and Loch, C. H. (2001). Project management under risk: using the real options approach to evaluate flexibility in R&D. *Management Science*, 47(1):85–101.

Lerner, J. and Nanda, R. (2020). Venture capital's role in financing innovation: what we know and how much we still need to learn. *Journal of Economic Perspectives*, 34(3):237–261.

Linton, J. D. (2016). Improving the peer review process: capturing more information and enabling high-risk/high-return research. *Research Policy*, 45(9):1936–1938.

Lotka, A. J. (1926). The frequency distribution of scientific productivity. *Journal of the Washington Academy of Science*, 16:317–323.

Mansfield, E., Rapoport, J., Romeo, A., Wagner, S., and Beardsley, G. (1977). Social and private rates of return from industrial innovations. *The Quarterly Journal of Economics*, 91(2):221–240.

Markowitz, H. M. (1952). Portfolio selection. *Journal of Finance*, 7:77–91.

Markowitz, H. M. (1959). *Portfolio selection: efficient diversification of investments*. Wiley, New York, NY.

Markowitz, H. M. (1991). Foundations of portfolio theory. *Journal of Finance*, 46(2):469– 477.

Mitton, T. and Vorkink, K. (2007). Equilibrium under-diversification and the preference for skewness. *The Review of Financial Studies*, 20(4):1255–1288.

Nicholson, J. M. and Ioannidis, J. P. A. (2012). Research grants: conform and be funded. *Nature*, 492:34–36.

Norton, E. and Tenenbaum, B. H. (1993). Specialization versus diversification as a venturecapital investment strategy. *Journal of Business Venturing*, 8(5):431–442.

Pleskac, T. J. and Hertwig, R. (2014). Ecologically rational choice and the structure of the environment. *Journal of Experimental Psychology: General*, 143(5):2000–2019.

Scherer, F. M. and Harhoff, D. (2000). Technology policy for a world of skew-distributed outcomes. *Research Policy*, 29(4-5):559–566.

Scott, R. C. and Horvath, P. A. (1980). On the direction of preference for moments of higher order than the variance. *The Journal of Finance*, 35(4):915–919.

Seglen, P. O. (1992). The skewness of science. *Journal of the Association for Information Science and Technology*, 43(9):628–638.

Simkowitz, M. A. and Beedles, W. L. (1978). Diversification in a three-moment world. *Journal of Financial and Quantitative Analysis*, 13(5):927–941.

Tobin, J. (1958). Liquidity preference as behavior towards risk. *Review of Economic Studies*, XXV(February):65–85.

Van Bekkum, S., Pennings, E., and Smit, H. (2009). A real options perspective on R&Dportfolio diversification. *Research Policy*, 38(7):1150–1158.

Veugelers, R. and Zachmann, G. (2020). Racing against COVID-19: a vaccines strategy for Europe. Policy Contribution n. 7, Bruegel.

Vilkkumaa, E., Liesiö, J., and Salo, A. (2014). Optimal strategies for selecting project portfolios using uncertain value estimates. *European Journal of Operational Research*, 233(3):772–783.

Viner, N., Powell, P., and Green, R. (2004). Institutionalized biases in the award of re- search grants: a preliminary analysis revisiting the principle of accumulative advantage. *Research Policy*, 33(3):443–454.

Von Neumann, J. and Morgenstern, O. (1953). *Theory of games and economic behavior*. Princeton University Press, Princeton, NJ.

Wang, D., Song, C., and Barabási, A.-L. (2013). Quantifying long-term scientific impact. *Science*, 342(6154):127–132.

Wang, J., Veugelers, R., and Stephan, P. (2017). Bias against novelty in science: a cautionary tale for users of bibliometric indicators. *Research Policy*, 46(8):1416–1436.

Yudkowsky, E. (2008). Cognitive biases potentially affecting judgment of global risks. In Bostrom, N. and Ćirković, M. M., editors, *Global catastrophic risks*, pages 91–119. Oxford University Press, New York, NY.