The relevance and the limits of the Arrow-Lind Theorem

Luc Baumstark

University of Lyon

Christian Gollier

Toulouse School of Economics

July 2013

1. Introduction

When an investment project yields socio-economic net benefits that are uncertain but independent of the systematic risk of the economy, these benefits should be discounted at the risk free rate if they are disseminated among a large population of stakeholders. This may be the case of a public project whose benefits are distributed within the large population of taxpayers. This is the essence of the Arrow-Lind Theorem, which played a crucial role in the evaluation of public policies around the world since its publication in 1970. Because of the differentiated treatment of investments evaluation that this result supports, it has also been extremely controversial. By reducing the discount rate to evaluate risky projects in the public sphere, it has certainly contributed to the expansion of the public sector in several western countries over the last four decades. We hereafter explain why this has been a mistake due to a fallacious interpretation of the Theorem.

1. The fundamental argument: Second-order risk aversion

In an economy of von Neumann-Morgenstern consumers with identical risk aversion and riskless income levels, it is optimal to maximize the dissemination of a given collective risk by a fair sharing of its burden in the population. The intuition of this result is based on the Arrow-Pratt approximation which states that the cost of risk is approximately proportional to its variance, or equivalently to the *square* of its size. This means that risk aversion is a second-order phenomenon (Segal and Spivak, 1990).[1] This implies that if each of the *n* agents bears *1/n* of the collective risk, each of them bears a cost of risk proportional to $1/n^2$, yielding a collective cost of risk proportional to $n/\ n^2=1/n$. When *n* tends to infinity, this risk dissemination washes out the collective cost of risk. This can be better understood in Figure 1, where we represented the quadratic relation between the size of individual risk of the individual cost of risk. When the collective risk is better disseminated, as when moving from A to B, the marginal cost of risk tends to zero, so that at the limit when $B \to 0$, this cost vanishes completely.

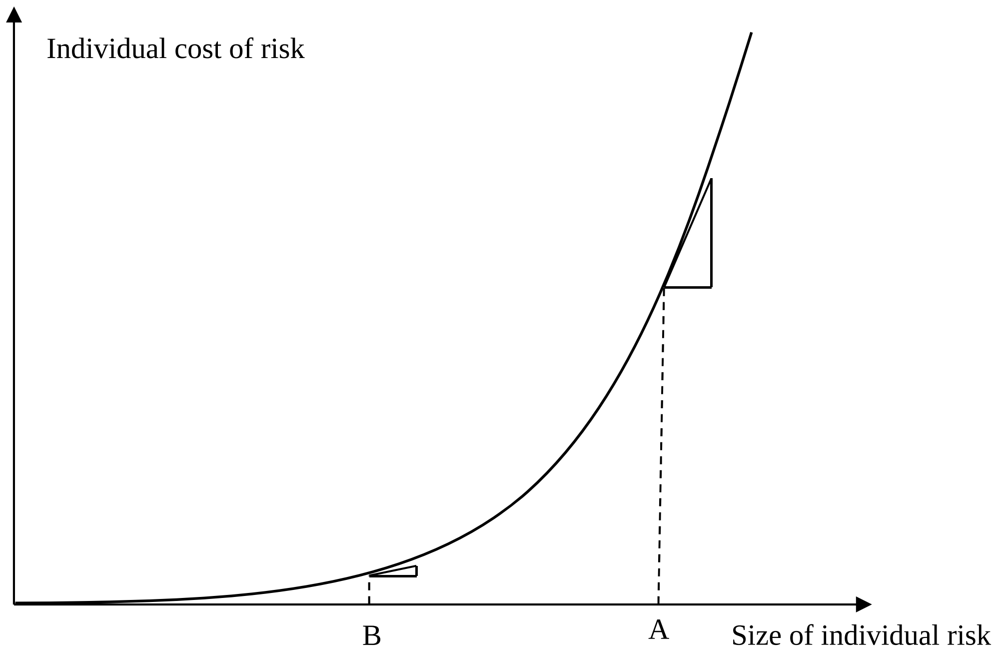

Figure 1: The individual cost curve of risk

This result is robust to the introduction of a "background risk" on the initial wealth levels of the stakeholders, as long as the risk involved in the collective project is independent of these background risks (Gollier and Pratt, 1996). The only effect of this background uncertainty is to raise the induced risk aversion to be used in the Arrow-Pratt approximation to evaluate the individual costs of risk.

---

[1] In many models alternative to expected utility, as the rank-dependent expected utility model (Quiggin, 1992), risk aversion is a first-order phenomenon and the Arrow-Lind Theorem does not hold in these contexts. But the expected utility theory and its associated independence axiom have a very strong normative appeal, so we don't attach much weight to this potential critique.

The publication of Arrow-Lind's paper was concomitant with the golden age of the Capital Asset Pricing Model (CAPM), and its improved version, the Consumption-based CAPM (CCAPM, Lucas, 1978). In the CCAPM as in the Arrow-Lind Theorem, idiosyncratic risks should not be priced in efficient financial markets. More generally, let us consider an investment project whose flow of net benefits $(F_0, F_1, F_2, ...)$ is statistically linked to the flow of aggregate consumption $(c_0, c_1, c_2, ...)$ in the economy through the following log-linear relation:

$$\ln F_t = \alpha + \beta \ln c_t + \varepsilon_t, \tag{1}$$

for all $t > 0$, where $(\varepsilon_1, \varepsilon_2, ...)$ is i.i.d. and normal. If the stochastic process for aggregate consumption is a geometric Brownian motion and if relative risk aversion is constant, then it is socially desirable to implement this project if and only if the Net Present Value (NPV) of its flow of expected net benefits is positive, with a discount rate being equal to

$$r = r_f + \beta\pi. \tag{2}$$

In this fundamental CCAPM formula, r_f is the risk free rate, π is the systematic risk premium, and β is the so-called beta of the project. This formula generalizes the Arrow-Lind Theorem by showing that when the project is not correlated to the systematic risk in the economy ($\beta = 0$), then the project should be discounted at the risk free rate. However, when the correlation between F_t and c_t is positive ($\beta > 0$), then the discount rate to be used to evaluate this project should be larger than the risk free rate. The intuition of this result is also based on the Arrow-Pratt approximation: When the risk of the project has a statistical component aligned on the systematic risk of consumption, this means that the representative agent in the economy already bears some positive component of it. In other words, contrary to lying at the origin of the axes in Figure 1 as assumed in the Arrow-Lind Theorem, the representative agent is positioned at some point A originally. Said differently, the second-order risk aversion argument underlying the Arrow-Lind Theorem does not apply here, because the risk is not marginal even after its dissemination. In this context, a perfect dissemination of the risk of the project does not eliminate the collective cost of risk, because the individual marginal cost of risk is not zero at point A. So, this increase in the collective risk that the project generates should be recognized in the evaluation process. This should be done by raising the discount rate. By how much depends upon the size of the systematic risk premium π and the projects' beta. Because the systematic risk premium on markets has been somewhere $\pi = 3\%$ and $\pi = 6\%$ over the last century (see for example Gollier 2012), to be compared to a risk free rate around $r_f = 0.5\%$, the effect of risk on the discount rate should be a crucial element of the evaluation process.

2. *Critiques to the Arrow-Lind Theorem*

Although, Arrow and Lind recognize in their paper that their result holds technically only for idiosyncratic risks, they support the idea that it has a much broader domain of applications. They first claim that "*the government undertakes a wide range of public investments and it appears reasonable to assume that their returns are independent*." In other words, they suggest that the average project in the economy should have a zero beta. This cannot be true. Because the economy can be represented as a portfolio of projects, the mean beta should be 1, and there is no reason to believe that the public sector has a portfolio of projects whose betas are systematically downward biased. Quite to the contrary, many public projects have large betas. Let me illustrate this with two examples. The infrastructure of highways is often justified on the basis of the time gained by their users, and by the number of lives saved. But the elasticities of the value of time and of the value of life with respect to changes in GDP are often assumed to be large. This implies a large beta for highway projects. In the case of the construction of a new high voltage line of electricity transportation justified by the anticipated increase in demand for electricity in a specific isolated region, the value of the project is positive only if the regional economy will indeed be growing in the future. This also yields a large beta. The same argument applies for fast train lines.

Arrow and Lind (1970) also justify the recommendation to discount public projects at the risk free rate r_f by suggesting that the systematic risk premium should be close to zero: "*It is sometimes argued that the returns from public investments are highly correlated with […] the business cycle. However, if we assume that stabilization policies are successful, then this difficulty does not arise.[…]* ". The facts have clearly contradicted this claim over the last four decades, and the systematic risk premium remained large on financial markets.

The bottom line has however been that the Arrow-Lind Theorem has often been interpreted in its broader, fallacious, sense. It is hard to evaluate the consequence of this wrong interpretation, but our own experience in France and elsewhere tells us that it has been sizeable. In the 80's and 90's, many French public firms have justified their investments on the basis of a low discount rate on the basis of the Arrow-Lind Theorem. For example, it is likely that the French nuclear industry could not have attained its full scale (75% of French electricity is produced using the nuclear technology) without the Theorem, before the privatization of Electricité de France 10 years ago (which indeed radically transformed the way in which EDF evaluated its investment projects). We also believe that many public-private partnerships exist just because of the discrepancies in the way the two sectors evaluate the cost of risk. This potentially generated a massive transfer of risk from the private sector to the public one.

In the academic world, the influence of the Arrow-Lind Theorem remains strong, but is somewhat more implicit. For example, in the context of climate change, most of the debate on the climate discount rate and on the social cost of carbon that followed the publication of the Stern Report (2007) relied on the Ramsey rule, which characterizes the socially efficient discount rate r_f to be applied to *safe* projects (see Gollier, 2012 for a survey). This approach is correct only if we believe that the benefits of fighting climate change for future generations are not correlated with the level of development that they will achieve. This is quite unrealistic. In particular, a large growth of consumption will generate more emissions and therefore larger marginal climate damages. Thus, it is likely that the climate beta is positive, so that the Arrow-Lind Theorem cannot be applied for climate change.

Arrow and Lind convincingly argue that the dissemination of risk and the risk sharing are not efficiently organized by financial markets, so that their result cannot be applied to the private sector. Some insurance markets are missing, whereas others are plagued with adverse selection and moral hazard problems. This implies that the CCAPM formula (2), which relies on efficient risk sharing in the economy, can only be interpreted as a rule-of-thumb for the evaluation of private projects. Investors and firms should take into account the fact that some of the risks generated by the project are retained by a limited number of stakeholders, which implies a positive collective cost of risk even when it has a zero beta. But it should also be noticed that the public sector also face frictions and inefficiencies from the same diseases. Various principal-agent problems force States to limit the dissemination of risk in the economy, as shown for example by Laffont and Tirole (1993). The goals of public servants are rarely aligned with the general interest, so that some risky rent should be allocated to them in order to provide better incentives. There is no reason a *priori* to believe that the public sector is more efficient than the private sector to disseminate risk in the population. In fact, there are some reasons to believe that the opposite is true. For example, financial markets are in a better position than states to disseminate country-specific risks around the planet.

3. *Concluding remarks*

In spite of its limited domain of applicability, the Arrow-Lind Theorem has probably played a crucial role in the development of the public sector in many countries over the last 40 years. Relying on a wrong interpretation of the Theorem, some lobbies have used this result to support their investment projects whose expected rate of return was not large enough to compensate for the increased systematic risk that they imposed on their stakeholders. Many countries have reacted to this behavior by imposing a large discount rate in the public sector. This is a second-best strategy, which implies not enough public investments in safe projects, and too many investments in risky projects. The first-best solution is to use risk-sensitive discount rates following the CCAPM rule (2), as has been done in Norway and in France over the last few years (see Gollier, 2011). Because of the well-known puzzles in finance, there is no global agreement about the levels of the risk-free rate r_f and of the systematic risk premium π to calibrate this formula. For long maturities, we believe that $r_f = 2\%$ and $\pi = 3\%$, so that socially efficient discount rates *r* should be set around $(2+3\beta)\%$, where β is the beta of the project under consideration.

References

Arrow, K.J., and R.C. Lind, (1970), Uncertainty and the evaluation of public investment decision, *American Economic Review*, 60, 364-378.

Gollier, C., (2011), *Le calcul du risque dans les investissements publics*, Centre d'Analyse Stratégique, Rapports & Documents n°36, La Documentation Française.

Gollier, C., (2012), *Pricing the planet's future: The economics of discounting in an uncertain world*, Princeton University Press, October 2012.

Gollier, C. and J.W. Pratt, (1996), Risk vulnerability and the tempering effect of background risk, *Econometrica*, 64, 1109-1124.

Laffont, J.-J., and J. Tirole, (1993), *The theory of incentives in procurement and regulation*, The MIT Press, Cambridge, MA.

Lucas, R., (1978), Asset prices in an exchange economy, *Econometrica*, 46, 1429-46.

Quiggin, J., (1982), A theory of anticipated utility, *Journal of Economic Behavior and Organization*, 3, 323-43.

Segal, U. and A. Spivak, (1990), First order versus second order risk aversion, *Journal of Economic Theory*, 51, 111-125.

Stern, N., (2007), *The Economics of Climate Change: The Stern Review*, Cambridge University Press, Cambridge.