

Risk Management as a Tool for Sustainability

Frank C. Krysiak

ABSTRACT. Although risk and uncertainty are inevitable aspects of the sustainability problem, they are often neglected in the sustainability discourse, especially in the economic analysis of sustainable development. We argue that this deprives the sustainability discourse of interesting connections to risk management. We show that defining sustainability as the obligation to limit the risk of harming future individuals provides a framework in which tools from risk management, like mean-variance analysis, can be employed to analyze planning decisions and to calculate a risk-minimizing policy mix. Furthermore, we discuss whether such a notion of sustainability can be an ethically tenable sustainability concept and how a positive probability of harming future individuals might be defended.

KEY WORDS: altruism, fairness, mean-variance analysis, risk, sustainability, uncertainty

Introduction

Sustainability has become a major framework for evaluating decisions with long-lasting consequences both on a political and on a firm-level scale. However, although much progress has been made in defining, measuring, and applying sustainability, some fundamental problems have received surprisingly little attention. One of these problems is the relation between risk and sustainability.

In a perspective slightly broader than the Brundtland definition, sustainability is a framework for assessing the impact of present decisions on the situation of future individuals. Any definition of sustainability that is covered by this perspective has to consider the future consequences of present actions. Any non-dictatorial notion of sustainability has also to be sensitive to future individuals' conceptions of well-being. Thus, at the core of sustainability lies futurity.

However, a necessary implication of futurity is uncertainty. Neither can we possibly know what conceptions of well-being future individuals will have as they are not in existence yet. Nor can we predict the future consequences of present decisions with certainty. All we can do is to assess the risks that present decisions impose on future individuals. So sustainability needs to be stated in terms of risks rather than in terms of certitudes.

Surprisingly, formal models of sustainability, especially in economics, have largely ignored this point. Either the uncertainty of the consequences of present decisions is neglected or it is taken as a reason for using vague definitions of sustainability. While these approaches have some merits in the political arena, they are highly problematic whenever sustainability is used in the context of a decision process. Neglecting uncertainty is problematic, because it allocates all risk-related costs to future generations, which contradicts the main tenet of sustainable development. Vague definitions of sustainability may be helpful to gain political acceptance (cf. Robinson, 2004). However, usually they do not result in operational concepts that can be used for choosing between different actions.

So at least in a decision context, ignoring risks is not a viable option. Indeed, outside the sustainability discourse, decisions with long-lasting consequences are frequently analyzed in a stochastic framework. Decision theory provides an extensive set of tools for such an analysis. However, these tools stem from a purely positive context. They can be used to analyze how a given individual's interest is best furthered in a risky world. However, they shed only a dim light on the question of how risks should be allocated among different individuals.

In this article, we use a recent concept of sustainability from economic literature and show that it

provides a link between an ethically tenable notion of sustainability and often-used tools of risk management. This concept is based on limiting the probability that a future generation is harmed, where a fairness-based criterion is employed to decide what constitutes “harm.” We show that it establishes an interesting connection with some tools of risk management. This connection permits adoption of these tools for assessing the sustainability of actions whose consequences are uncertain and facilitates finding a policy mix that minimizes the risk of harming future individuals.

Our analysis is based on a decision theory framework that can be used on any scale. It applies to a firm interested in analyzing the sustainability of different actions, to a regional administration inquiring about the sustainability of development plans, as well as at a national and international level, where sustainability relates to nearly all political decisions.

In the next section, we set up our framework and highlight the connections between sustainability and risk management. In section [The ethical dimension of a risk-based concept of sustainability](#), we discuss possible ways to construct an ethical foundation for a risk-based notion of sustainability. Section [Discussion and conclusions](#) concludes.

Sustainability as risk containment

If we neglect studies that make only passing references to risk and uncertainty without adapting their analysis to the specific problems caused by uncertain outcomes and uncertain preferences, formal concepts of sustainability that account for uncertainty are rare in the literature. Some studies like Woodward (2000) and Asheim and Brekke (2002) consider outcome uncertainty; that is, unpredictable influences on the future consequences of present decisions. Other studies, like Heal et al. (1998), Ayong Le Kama (2001), Ayong Le Kama and Schubert (2004), and Krysiak and Krysiak (2006), inquire about the effects of preference uncertainty; that is, the imperfect predictability of future individuals’ evaluations of their situation.

In this article, we use the criterion of sustainability advanced in Krysiak and Krysiak (2006),¹ which, in turn, is based on combining the fairness-based

definition of sustainability given in Woodward (2000) with the theory of social targeting developed in Naga (2003), Cornia and Stewart (1995), and Sen (1995).

This criterion holds that an action is sustainable if the probability that a future generation is harmed by this action is limited by a constant α . A future generation is seen as being harmed, whenever it would have preferred the state of the world without the action to that with the action.² This criterion contains some often discussed concepts of sustainability as special cases; that is, as special choices of α . As shown in Krysiak and Krysiak (2006), it includes weak sustainability for $\alpha = 50\%$ and strong sustainability for $\alpha = 0\%$, if some assumptions on preferences, technological possibilities, and expectations are met. For any $\alpha < 100\%$, the criterion can also be seen as a version of non-declining welfare.

This definition of sustainability is based on a notion of fairness that is defined relative to a *status quo*; that is, to the situation unchanged by actions of the present generation. Future generations have a “right of veto” against actions that they regard as being detrimental to their situation. However, they have no legitimate claim to demand changes to the *status quo*. The definition is based on fairness in that it employs an envy-based criterion³ for deciding which actions are potentially harmful to future individuals. Given that envy is based on the comparison of states from the perspective of a single individual, the above sustainability criterion is free of interpersonal welfare comparisons. It can thus be used even if present and future generations do not share a common notion of well-being.

The relation between risk and sustainability

The above definition of sustainability provides a useful framework to discuss the relation between risk and sustainability. Consider a decision between several actions that have long-lasting consequences. This decision may be a firm’s decision between investment options, a planning decision of a regional authority, or a political decision at national or international levels.

For each action, we could calculate the expected impact on the well-being of every future generation. However, due to the uncertainty of future

preferences and the imperfect predictability of the actions' outcomes, the actual consequences will usually deviate from these expected impacts. Planning with only the expected impacts would allocate the risk of such deviations completely to future generations, which is not consistent with the notion of sustainability. So for assessing the sustainability of an action, we need to include information about the possible deviations from the expected impacts. The above concept of sustainability uses such information to calculate and limit the probability of harming a future generation.

Let us consider a fictive example to highlight this approach. Assume that the use of a given area of land has to be decided. There are three alternative proposals: The land could be used as a recreational area, it could be zoned for residential use, or it could be used for industrial production. Assume that residential and industrial use have a higher expected benefit than recreational use, but also share increased risk; the benefit and risk of industrial use exceeding that of residential use. In all three cases, the uncertainty is a conglomerate of outcome uncertainty (e.g., it is not sure which social benefits the industrial use will have) and preference uncertainty (e.g., future preferences for recreation are uncertain). For presentational simplicity, we distinguish only between close and distant future generations.

In Table I, we have set fictive values for the expected benefits (relative to the *status quo*) and the standard deviation of the benefits for each alternative. In practice, establishing such numbers will be difficult at best. However, although this approach seems to be complicated, variants of it are commonly used for informing political decisions. For example, the IPCC reports on climate change contain different scenarios for global warming, each of which is

given some credence, which might be used to construct expected values and risk measures.

To evaluate this information according to the above criterion of sustainability, we identify for each possible action and for each future generation the probability that the action results in a situation that is (from the view of the generation in question) worse than the *status quo*. For each action we then identify the future generation that is most likely to be harmed and use the probability of harming this generation as the rating of this action.

In this way, we can rank the alternatives: the actions with a lower probability of harming a future generation are better according to the above criterion than those with a higher probability of doing harm. Furthermore, we can identify which alternatives remain, if we wish to constrain the probability of harming a future generation to some level α .

For the data of Table I, we have depicted this approach in Figure 1.⁴ As this figure shows, this concept yields results that differ substantially from those based on a deterministic analysis. The alternative "industrial use" has the highest expected value and would thus be chosen in a deterministic context. However, due to the higher variance of its outcomes, it has a greater probability of harming the depicted generation than the other alternatives. Indeed, it is the worst choice, in the sense, that it results in the highest risk of being harmful to future generations. The alternative "residential use" yields a better balance between expected benefit and risk. It implies the lowest probability of harming a future generation and would thus be recommended by our approach, if a single alternative has to be chosen.

As this simple example shows, it is essential to consider not only the expected consequences when assessing the sustainability of projects but also the

TABLE I

The expected benefits (μ) relative to the *status quo*, and the standard deviation of benefits (σ) for close and distant generations for the three land use decisions (fictive example)

Usage	Close generation		Distant generation	
	μ	σ	μ	σ
Recreational	0.33	0.1	0.33	0.25
Residential	0.66	0.2	0.66	0.4
Industrial	1	0.5	1	1




Figure 1. The distribution (under the normal distribution assumption) of the evaluated outcomes x for the three land use alternatives. The right plot is a magnified depiction of the rectangular area indicated in the left plot. The vertical line in the right plot represents the *status quo*. The shaded areas depict the probabilities of harming the generation for each alternative.




Figure 2. Lines of constant probabilities of harming a future generation for a normal distribution depicted in dependency of the standard deviation of benefits σ and of the expected benefit μ (measured relative to the *status quo*). The circles correspond to the outcome of the three alternatives of Table I for close generations and the dots to those for distant generations.

projects' risks. The concept outlined above provides a framework for analyzing the trade-off between the expected benefit a project might have for future generations and the risk that it imposes upon them.

In Figure 2, we have depicted this trade-off for the case of a normal distribution of future benefits. The lines in the figure represent all combinations of expected benefit and standard deviation that imply the same probabilities of causing harm. The nearer the line approaches the lower right corner of the diagram, the smaller is this probability, and, consequently, the notion of sustainability that is associated with it will pose greater demands (i.e., it corresponds to a lower α).

With such a diagram, we can evaluate action proposals according to their sustainability. In the

example depicted in Figure 2, we see that it is always the “distant generation” that is most likely to suffer. With respect to this generation, the alternatives “recreational use” and “residential use” are sustainable at the level $\alpha = 10\%$. However, the alternative “residential use” is preferable to the alternative “recreational use,” because the former is even sustainable at the level $\alpha = 5\%$, whereas the latter is not. Thus, we can conclude that the alternative “residential use” is the preferable choice and that it imposes a probability of harming future generations that is, at most, 5%.

Employing risk management tools for sustainability

The above analysis has shown that it is possible to transfer some tools of decision theory to the sustainability discourse on the basis of a rigorous definition of sustainability. This approach has the additional advantage that we can not only use the descriptive techniques outlined above, but can also employ concepts from risk management to search for policies that actively reduce the risks that present actions impose on future generations.

For this, we need information about the interrelationship of consequences arising from different alternatives. In many settings, future constellations of preferences and uncertain influences on outcomes that render one alternative disastrous will render another alternative advantageous, and vice versa. Thus, the evaluated outcomes of different alternatives will often be correlated. Such correlations are usually measured by correlation coefficients like those displayed in Table II for the fictive example.

If such correlations exist, it will usually be advantageous to implement a policy mix instead of choosing a single alternative. By using a portfolio of alternatives, for example, a diversified energy system, the risk of disastrous outcomes can be decreased. The question is how a reasonable policy mix can be calculated.

Modern portfolio theory enables us to address this question. The above approach is able to adopt methodological applications from this literature because it formulates the sustainability problem as a decision between risk and expected benefit. The literature provides an extensive set of tools for this purpose. However, for presentational simplicity, we

TABLE II

Correlation coefficients between the risks for distant generations for the three land use decisions (fictive example)

Usage	Recreational	Residential	Industrial
Recreational	1	-0.5	0.4
Residential	-0.5	1	-0.6
Industrial	0.4	-0.6	1

use only the simplest possible construct that is applicable in our setting. This is mean-variance analysis, which dates back to Markowitz (1952) and which, although being largely discarded in economic theory,⁵ is still frequently used in applied finance.

Mean-variance is focused on calculating combinations of alternatives that result in the smallest variance of returns for a given expected return. In our context, we can use this approach to calculate a risk-minimizing portfolio of actions.

Let us perform such an analysis for the land use example by assuming that we can implement not only each alternative, but also all combinations of the three alternatives; for example, by using a part of the land for each alternative. The question is how much land should be reserved for which usage.

Following the mean-variance concept,⁶ we can calculate the efficient frontier. The efficient frontier describes the minimal risk for a given expected benefit that is achievable by some mix of the alternatives. In Figure 3, we have depicted this minimal risk as a function of the expected benefit.

This figure indicates that, except for the points on the boundary (i.e., an expected benefit of 1/3 or 1), it is always recommendable to use a policy mix: for all points on the efficient frontier, at least two land uses have a share of the land that is strictly greater than zero.




Figure 3. The efficient frontier calculated from a mean-variance approach with the data of Tables I and II (Plot a), and the shares η of the three alternatives on the efficient frontier (Plot b).

This is a standard result. However, it is interesting to combine it with the above risk-based perception of sustainability. In Figure 4, we have simultaneously depicted the efficient frontier of Part (a) of Figure 3 with the levels of sustainability shown in Figure 2.

There are two ways to use this figure. First, we can set a maximal probability of harming a future generation and accept all alternatives that achieve this goal. In this case, we use the corresponding thin line in Figure 4 and all points on the efficient frontier (thick line) that lie on or below this thin line represent acceptable efficient alternatives.

Second, we can try to minimize the probability of harming a future generation. This is achieved at the (usually unique) point at which a thin line is tangential to the efficient frontier. In our setting, this point is given by $\mu \approx 0.47$, $\sigma \approx 0.156$, which according to Figure 3b corresponds to 59.55% of the land being dedicated to recreational use, 40.24% to residential use, and 0.21% being reserved for industrial use.




Figure 4. The efficient frontier (thick line) from Figure 3 (Plot a) is combined with the lines of constant probabilities of harming a future generation from Figure 2 (thin lines). The thin lines correspond to probabilities of harming a future generation of 20%, 10%, 5%, 1%, 0.5%, 0.132%, and 0.01% (from the upper-left to the lower-right line).

With this combination of alternatives, the risk of harming a future generation is reduced to 0.13%.

Comparing this to the case where only a single alternative is used demonstrates the potential power of the application of risk management techniques to the problem of sustainability: the best single-alternative choice (residential use) implies a risk of 4.779% of harming a future generation; the optimal policy mix reduces this to 0.13%. If we had neglected the risk altogether, the likely choice would have been industrial use. This decision would have harmed a future generation with a probability of nearly 16%, that is, more than a hundred times the minimal achievable risk.

Of course, this is only a fictive example. However, as mean-variance analysis shows, the main qualitative result, that diversification can be used to reduce risks and thus to achieve a more sustainable outcome, is robust. The analysis could be redone with almost all reasonable values for the parameters given in Tables I and II and, although the resulting numbers will vary, the main message that a policy mix furthers sustainability will remain.

Indeed, this risk reduction effect is commonly taken advantage of on financial markets, where most investors hold a portfolio of assets to reduce the total risk by means of diversification. Our analysis shows that concepts that have proven to be valuable in finance can be applied to the analysis of sustainability. The necessary transfer of methods is not undertaken on an *ad-hoc* basis, but on the foundation of a definition of sustainability that includes several commonly used sustainability concepts.

This approach has the advantage that it allows not only for a rigorous application of the tools of risk management, but also for a detailed assessment of its ethical implications, as we show in the following section.

The ethical dimension of a risk-based concept of sustainability

Until now, we have sidestepped the question of whether limiting the probability of harming future generations is an ethically tenable concept of sustainability. Krysiak and Krysiak (2006), from whom we have adapted this concept, do not address this point in depth. Furthermore, even if we can give a

tentatively positive answer to this question, we still have to ask how an “acceptable” risk of harming future generations can be chosen. We discuss these points successively.

Is risk containment an ethically tenable approach to sustainability?

The first important observation is that in a world where the consequences of today’s actions can only be imperfectly predicted and where we cannot possibly know what state of the world future individuals will prefer, there is only a limited chance of certainly not harming future generations. Furthermore, as Krysiak and Krysiak (2006) have shown, such a demanding protection of future individuals (which corresponds to setting $\alpha = 0\%$) will severely restrict our ability to better the fate of future people. Indeed, as they argue, the situation is somewhat similar to hypothesis testing in statistics: the more certain we like to be of not accepting a wrong hypothesis, the higher becomes the probability of rejecting a true hypothesis. The surer we wish to be of not harming future generations, the more likely we will decline actions that would benefit both the present and future generations.⁷

The above definition of sustainability permits us to balance the risks of harming and not benefiting future individuals. Given that actively harming someone is commonly seen as worse than passively not bettering someone’s situation (even if both decisions result in the same deprivation),⁸ the former risk should presumably be lower than the latter, which in our context requires us to set a suitably low level for α . However, in most cases, it will be reasonable to allow for some positive risk of harming future individuals, because, otherwise, the possibilities for raising the welfare of future and present individuals would be severely reduced, often leaving only the option of preserving the *status quo*. Indeed, as we will argue below, given only information that is currently available, future individuals would insist on taking on some risk in order to achieve a higher expected standard of living. Thus, there is a level $\alpha > 0$ that is acceptable for all future individuals.

So the basic concept of allowing for some probability of harming future individuals seems to be

tenable. However, in order to be acceptable, the concept must also reflect adequate congruence between risk-induced costs and the responsibility for the risk-causing action. Since we have two sources of uncertainty, this translates into two distinct questions. First, owing to preference uncertainty, we have the question of responsibility for preferences, which has led to the distinction between welfare- and resource-based concepts of distributional justice in Dworkin (1981a, b).⁹ Second, owing to outcome uncertainty, we have the question of who carries the risk of adverse outcomes of present actions.

Concerning the responsibility for preferences, the above approach – as a fairness-based construct – pursues a way in between welfare- and resource-based concepts of distributional justice. It is based on individually evaluated outcomes and thus on welfare. So the costs of preference heterogeneity and preference uncertainty are socialized. If we hold each generation responsible for its preferences, this implies a higher level of intergenerational risk sharing than Dworkin (1981b) proposes. However, due to being defined relative to the *status quo*, this sharing mechanism is constrained to impose no undue burden on the present generation. It grants claims for more resources that are based on uncommon preferences only if these preferences could have been better met in the *status quo*. So the present generation is held responsible for paying regard to future preferences when effecting changes to the *status quo*. However, the present generation does not have to change the *status quo* in a way that is detrimental to itself. In other words, each generation needs to use only the additional welfare that it generates in changing the world in order to compensate future generations for possible hardships resulting from extravagant preferences. Thus, the concept seems to provide a reasonable balance between intergenerational risk sharing and individual responsibility for preferences.

Concerning outcome uncertainty, each generation is responsible for bearing most of the risk that it causes by changing the *status quo*. In our concept, future generations have a “right of veto” against actions that they consider to be detrimental to their well-being. Thus, if the present generation wants to implement changes to the *status quo* that impose some risk on future generations, this is only possible if these risks are accompanied by an increase in

expected benefits that all future generations consider as being an adequate compensation.

Altogether, the above concept of sustainability seems to achieve at least some congruence of responsibility and cost-bearing. Concerning preference uncertainty, future generations can claim compensation only if changes to the *status quo* have worsened their situation. They are thus, in line with Dworkin (1981b), considered to be largely responsible for their preferences. Concerning outcome uncertainty, each generation bears only the (risk-induced) costs of those actions from which it potentially profits.

Choosing an “appropriate” risk of harming a future generation

Let us now turn to the question of how an “acceptable” probability of harming a future generation (i.e., α) can be chosen. There are two main possibilities. If we consider sustainability to be a descriptive concept, then the value of α is simply a measure of altruism, which can be observed but not judged.

In contrast, if we perceive sustainability as a normative concept, then α has ethical implications. In this context, it is informative to construct sustainability by a contractarian exercise along the lines of Rawls (1971) and Gauthier (1986). Note that in contrast to our analysis in section [Sustainability as risk containment](#), the rational choice framework is essential for this argument.

As argued by Sugden (1993), this enterprise can only succeed if we introduce some moral constraints *a priori*. For this, imagine that the present and all future generations come together behind a specific variant of Rawl’s veil of ignorance. Behind this veil, each generation knows when it will live but has only the information concerning future preferences and future outcomes that the present generation has outside the veil. Furthermore, behind this veil, the decision for an action shall require unanimity.

In such a setting, each generation will decide upon its future fate with the information available at present. Thus, each generation is subject to the trade-off between increasing its expected benefit and the accompanying change in its risk. Based on its own risk

aversion, it will choose a level of α . Whenever there is no infinite risk aversion and whenever an increase in expected benefits has to be bought with an increase in risk,¹⁰ this α will be strictly greater than zero. Due to the unanimous decision rule behind the veil, the generation that is most likely harmed will determine the final choice of an action.

In this way, an action is agreed upon, if for each generation, the probability of being harmed is below the generation's chosen value of α . So α can be constructed from the risk preferences of the future individuals. Indeed, this construction closely resembles the Rawlsian defense of a maximin decision rule in social choice theory (cf. Roemer, 1996, ch. 5).

This exercise shows that a strictly positive probability of harming future generations can be defended, because given the presently available information, future generations will demand such a setting of α . Furthermore, this argument suggests that in a rational choice framework, α should be related to risk aversion. Thus, although we have to allow for possible differences between present and future risk preferences, currently observed levels of risk aversion may provide useful information for choosing an acceptable boundary to the probability of harming future generations.

Discussion and conclusions

In this article, we have proposed to use the fairness-based criterion of sustainability under uncertainty advanced in Krysiak and Krysiak (2006) in a broader context. We have shown that this criterion provides a link between commonly used tools of risk management and a potentially operational notion of sustainability. This link can be used to transfer risk management tools from their positive context to the normative setting in which questions of sustainability are commonly argued. Furthermore, we have addressed the question whether this definition of sustainability is ethically tenable and in which way its basic notion of an "acceptable" risk for future generations can be defended.

It need not be pointed out that this analysis does not provide a fully operational concept and that some of our arguments have a somewhat speculative touch. Furthermore, we have focused in our examples on a rather limited conception of sustainability by

excluding the social and cultural dimensions, and have employed only the simplest possible tools of risk management. Finally, we have constrained our exposition to the rational choice framework that is predominant in economics.

The latter three restrictions are mainly for presentational simplicity. The concept can handle any consequences of present actions, social, and cultural consequences included because it is based on the welfare of future individuals as they perceive it. The concept of welfare, as it is understood in economics, comprises everything that a person might regard as being relevant to evaluating his or her situation and is thus sufficient to capture a broad perspective of sustainability. Furthermore, since we have based our concept on a generation, we can also cover all forms of cultural and social interactions between individuals as long as these individuals perceive them as being relevant for their personal welfare.

Concerning the second restriction, the concept can be used in conjunction with more advanced tools of risk management, even with full-scale optimization techniques. However, although mean-variance analysis has been abandoned in economic theory, it is still a frequently used tool, and it has the advantage that it is easily communicable. Thus, it serves well for a simple exposition of the feasibility of connecting sustainability and risk management.

The restriction to the rational choice model is mainly for presentational purposes. With the exception of the discussion in section [The ethical dimension of a risk-based concept of sustainability](#), our analysis only requires that we can discern what harms a future individual. We have used individual preferences to define harm via an envy-based criterion. However, since we do not need a quantitative measure of harm, we could also use Sen's concept of capabilities or a resource-based concept, like the availability of basic goods. Thus, our main concept is applicable with constructs of well-being that differ substantially from those embedded in the rational choice framework.

Our approach could also be criticized for being overly preoccupied with risk as well as for using a rather specific approach to handling risks. Concerning the first point, we have argued that uncertainty is a necessary consequence of evaluating the future consequences of present actions. Thus, risk lies at the core of the sustainability problem.

Neglecting this point would defer the risks that result from present actions completely to future generations, which seems to be incompatible with any ethically meaningful concept of sustainability. So while it may be helpful to construct sustainability as a concept of weighing the fulfillment of present and future needs against each other in a general discourse, any application of sustainability in actual decision processes has to recognize that sustainability is essentially a weighing of risks and not of certitudes.

Concerning the second point, it is obvious that restricting the probability of unfavorable outcomes is a somewhat peculiar approach to handling uncertainty. Economists will argue that it is even an irrational approach, at least when considered in the framework of expected utility theory, and that using expected future welfare would be more consistent with economic theory. However, as Chichilnisky (2000) has shown, expected utility theory can be insensitive to catastrophic risks and is, thus, at least a debatable framework for discussing sustainability. Furthermore, the approach upon which our analysis builds is widely used in statistical test theory and, to a lesser extent, in the theory of social targeting. Finally, it has the advantage that it encompasses several often used concepts of sustainability and is consequently rather an extension of existing definitions than a new definition. Thus, although it is undeniably an unusual approach, we would argue that it is a particularly advantageous approach.

Notes

¹ Thereby, we proceed in the framework of rational choice; that is, we presume that individuals have rational preferences that are directly related to individual well-being. There exists an extensive body of literature that criticizes this framework and that provides possible alternatives (cf. e.g., Sen, 1982, 1999; Roemer, 1996). However, discussing the relative merits of the numerous alternative frameworks is beyond the scope of this article. Furthermore, we use the rational choice framework mainly to highlight the connection to economic decision theory. The analysis in this section requires only that we are able to decide whether an outcome of a present action will harm a future individual or not. To this end, we could use most of the alternative frameworks available, like capability- or resource-based approaches.

² In this article, we will not emphasize the distinction between harming future generations and harming

future individuals because this is not relevant to our analysis. In our context, a future generation can be seen as being harmed, whenever at least a fraction β of its individuals are harmed, with β taking any value between 0% and 100%. e.g., We could use unanimity or a simple majority rule.

³ In the definition commonly used in social choice, fairness consists of envy-freeness and efficiency, see Foley (1967) and Varian (1975). In this context, envy refers to a situation in which one individual prefers the situation of another individual to its own situation.

⁴ For this, we have further assumed that future benefits are normally distributed.

⁵ As Rothschild and Stiglitz (1970) have shown, mean-variance analysis is not generally consistent with expected utility maximization and thus results in inferior outcomes for most utility functions than more complex approaches, like full-scale optimization techniques.

⁶ For a detailed exposition see Markowitz and Todd (1987).

⁷ Note that this comparison refers only to the mathematical properties of hypothesis testing and the risk-based concept of sustainability. Test theory proceeds in a positive context, whereas sustainability is usually discussed in a normative context. Therefore, these approaches cannot be compared in general.

⁸ This is coded into law in most countries. The punishment for causing death by declining to help is usually substantially lower than that for actively contributing to someone's untimely end.

⁹ This distinction is mirrored in the sustainability discourse, although the link to responsibility for preferences is sometimes obscured by the common assumption of homogeneous preferences.

¹⁰ Given efficient behavior of all generations, the latter condition will always hold.

Acknowledgments

We would like to thank the editors of this special issue and two anonymous reviewers for their helpful comments.

References

Asheim, G. B. and K. A. Brekke: 2002, 'Sustainability When Capital Management has Stochastic Consequences', *Social Choice and Welfare* **19**, 921–940.

- Ayong Le Kama, A.: 2001, 'Preservation and Uncertain Future Preferences', *Economic Theory* **8**, 745–752.
- Ayong Le Kama, A. and K. Schubert: 2004, 'Growth, Environment and Uncertain Future Preferences', *Environmental and Resource Economics* **28**, 31–53.
- Chichilnisky, G.: 2000, 'An Axiomatic Approach to Choice under Uncertainty with Catastrophic Risks', *Resource and Energy Economics* **22**, 221–231.
- Cornia, G. A. and F. Stewart: 1995, 'Two Errors of Targeting', in D. van de Walle and K. Nead (eds.), *Public Spending and the Poor: Theory and Evidence* (John Hopkins University Press, Baltimore), pp. 350–386.
- Dworkin, R.: 1981a, 'What is Equality? Part 1: Equality of Welfare', *Philosophy & Public Affairs* **10**, 185–246.
- Dworkin, R.: 1981b, 'What is Equality? Part 2: Equality of Resources', *Philosophy & Public Affairs* **10**, 283–345.
- Foley, D.: 1967, 'Resource Allocation and the Public Sector', *Yale Economic Essays* **7**, 45–98.
- Gauthier, D.: 1986, *Morals by Agreement* (Oxford University Press, Oxford, New York).
- Heal, G. M., G. Chichilnisky and A. Beltratti: 1998, 'Uncertain Future Preferences and Conservation', in G. Chichilnisky and A. Vercelli (eds.), *Sustainability: Dynamics and Uncertainty*. (Kluwer Academic Publishers, Dordrecht), pp. 257–276.
- Krysiak, F. C. and D. Krysiak: 2006, 'Sustainability with Uncertain Future Preferences', *Environmental and Resource Economics* **33**, 511–531.
- Markowitz, H. M.: 1952, 'Portfolio Selection', *The Journal of Finance* **7**, 77–91.
- Markowitz, H. M. and G. P. Todd: 1987, *Mean-Variance Analysis in Portfolio Choice and Capital Markets* (John Wiley & Sons, New York).
- Naga, R. A.: 2003, 'The Allocation of Benefits under Uncertainty: A Decision Theoretic Framework', *Economic Modelling* **20**, 873–893.
- Rawls, J.: 1971, *A Theory of Justice* (Harvard University Press, Cambridge, MA).
- Robinson, J.: 2004, 'Squaring the Circle? Some Thoughts on the Idea of Sustainable Development', *Ecological Economics* **48**, 369–384.
- Roemer, J. E.: 1996, *Theories of Distributive Justice* (Harvard University Press, Cambridge, MA).
- Rothschild, M. and J. E. Stiglitz: 1970, 'Increasing Risk. 1. Definition', *Journal of Economic Theory* **2**, 225–243.
- Sen, A. K.: 1982, *Choice, Welfare, and Measurement* (Harvard University Press, Cambridge, MA).
- Sen, A. K.: 1995, 'The Political Economy of Targeting', in D. van de Walle and K. Nead (eds.), *Public Spending and the Poor: Theory and Evidence* (John Hopkins University Press, Baltimore), pp. 11–24.
- Sen, A. K.: 1999, *Commodities and Capabilities* (Oxford University Press, Oxford, New York).
- Sugden, R.: 1993, 'The Contractarian Enterprise', in D. Gauthier and R. Sugden (eds.), *Rationality, Justice and the Social Contract*. (Harvester Wheatsheaf, New York), pp. 1–23.
- Varian, H. R.: 1975, 'Distributive Justice, Welfare Economics, and the Theory of Fairness', *Philosophy & Public Affairs* **10**, 223–247.
- Woodward, R. T.: 2000, 'Sustainability as Intergenerational Fairness: Efficiency, Uncertainty, and Numerical Methods', *American Journal of Agricultural Economics* **82**, 581–593.

Department of Economics,
University of Basel,
Peter Merian-Weg 6, 4002 Basel, Switzerland
E-mail: Frank.Krysiak@unibas.ch