Rerum Causae, 6 (2014), 89-102

Decision-making, Uncertainty and Precaution: a Case Study from Formal Ethics

Camilla Colombo

Abstract

A controversial decision making problem is handled with Rational Decision Theory tools. First, I build a convenient decision model, catching the relevant features of the original example, concerning the legitimacy of a genetic engineering intervention. In this peculiar decision context, I propose the employment of the decision criterion of maximin as a rational solution for the dilemma. I justify this criterion with two main arguments, I regard as underlying characteristics of the decision context: the agent's state of uncertainty and the responsibility requirements involved. The suggested conclusion is shown to be an instance of the so-called Precautionary Principle, often advocated for decision-making facing environmental and public health risks. The objections of vagueness, pessimism and strictness attached to the principle are addressed to the same features of uncertainty and responsibility which justify maximin.

I. Alice's case

In a hypothetical technologically advanced scenario, a genetic engineering intervention for height enhancement is made available to prospective parents. Fetus undergoing this treatment will experience a 20cm height increase. The intervention must be authorized by parents. Alice, like all prospective parents,

faces the decision whether or not to authorize the height enhancement. She balances the benefits and the disadvantages of an above-average height: e.g., Alice's child would have more chances as a basketball player but fall short as a professional jokey or dancer. Alice could also reckon other elements, widening the evaluation to pretty appearence, uneaseness and every possible attitude towards tallness.

Before modeling Alice's problem, I would go through some characteristics of the decision she deals with. First, her choice deeply affects the life of another person, her unborn child. Despite this, the child can't express any preference, taste or inclination when the choice is made. Nevertheless, these features are assumed to determine the desiderability of an above-average height. So, Alice can't easily choose the alternative which would favour her child the most. Using RDT tools, the child's attitudes can be thought as the *states of nature* of a decision matrix and the strategies available to the agent as the *actions*. The *consequences* of Alice's decision are all the ordered couples of states and actions.

Formally, let $A_i = (a_1, ..., a_n)$ the set of all the actions available to an agent i, and $S = (s_1, ..., s_k)$ the set of all the states of nature. The states of nature can be regarded as different life plans which can be chosen by Alice's child. In standard RDT, they are defined as mutually exclusive and jointly exhaustive. Then, the set $C = (c_1, ...c_n)$ of the consequences is the product of all $A \times S$. A consequence c_j can be defined as a pair (a_i, s_k). Let's notice that an action $a \in A$ can be identified with the list of consequences it brings about in different states of nature.

In the following sections, I will discuss two relevant properties of Alice's decision problem.

II. Uncertainty

Let's go one step further in the formalisation of Alice's decision problem. E. g., a possible state of nature could be s_1 = 'the child wants to become a bas-

ketball player and finds above average height pretty'; the available actions are authorise (a) or not authorise ($\neg a$) the intervention. We could then require that Alice's preference relation $\succeq_A$ is such that $c_1 = (s_1, a) \succeq_A c_2 = (s_1, \neg a)$. A reasonable decision criterion would then impose Alice to choose a over $\neg a$ when the realisation of the world is s_1, since a leads to the preferred consequence.

In standard decision theory, a problem, like the one described above, can be classified on the base of the epistemic status of the agent towards the states of nature. In Alice's case, her attitude towards the states is defined as uncertainty: she doesn't know which state of nature will occur, i.e., which preferences her child will display, neither she could attach an objective probability to each state. In the realm of uncertainty, I propose to define Alice's epistemic status as *complete ignorance* in the sense used by Luce and Raiffa[1]. She doesn't have the faintest idea whether her child would find above average height good looking or ugly: there is no reliable process available for predicting her child's tastes, preferences and career choice.

But that's not the whole story. I suggest that Alice isn't even in the position to build a state space which satisfies the requirements stated in Section 1. Obviously, Alice can't list all the possible combinations of her child's preferences, desires and attitudes, since they are potentially infinite. Thus, the state space of the problem can't be exhaustive in this respect. Nevertheless, it can be argued that any state space partition is arbitrary, and depends on the features the agent considers relevant for the problem in question. Then, Alice could build a state space like this:

- s_1 = 'basketball player and finds above average height pretty'
- s_2 = 'jokey and feels awkward about height'
- s^* = 'any other height-neutral profession and indifferent to above-average height'

[1] Luce and Raiffa, *Games and Decisions*

where these three states are mutually exclusive and jointly exhaustive (s_1 and s_2 clearly affects the consequences of a height enhancement intervention).

My point is that such a procedure is neither psychologically plausible neither intuitive for Alice, and consequently an adequate model should consider a different account of the states of nature. First, there isn't an univocal and uncontroversial way to build the state space, since more detailed versions of s_1, s_2 and s^* relevant for the height enhancement could be reckoned. It's more reasonable instead to think that Alice doesn't even need a complete set of states of nature to take her decision. The situation she faces is so complex that she can only figure out some possible scenarios on the grounds of the consequences, which she can determine in an unambiguous and clear way. These scenarios can be thought as some highly distinguishable and relevant states which could affect the desiderability of the treatment, like s_1 or s_2. Then, the states she needs to distinguish are those which bring about different consequences for the height enhancement intervention. In other words, given that she only considers some of the states of nature, it's not sound require that the state space must be complete (i.e., that the states must be jointly exhaustive).

My proposal is that a suitable model for Alice's decision problem weakens the requirements over the state space and focuses on the consequences (a_i, s_k) of the actions chosen in some specific, relevant and highly distinguishable states. In Section 4, this insight will be modelled assuming that Alice's preference relation satisfies the axiom of column duplication.

III. Responsibility

Analysing Alice's epistemic status, I suggested that the agent focuses more on the consequences of her actions than on the states of nature, since she can't distinguish and list all the possible attitudes and careers (and she doesn't need to do so). In this section, I deal with another reason why Alice would rather concentrate on the consequences in general and on a specific subset of outcomes

in particular.

Alice's decision will deeply affect her child's prospectives and success: the agent doesn't decide for herself, but for another person whose preference profile is still undefined. I assume now that Alice cares about her child and desires to favour them in her chosen life plan; otherwise, we can regard this property as a psychological interpretation of Alice's revealed preferences. In other words, Alice feels *responsible* for the consequences of her actions, and wishes she could justify her choice to her child. Under this assumption, her favourite outcome would be (s_1, a), since she would act in accordance with her child's preferences. On the other hand, she wants to avoid (s_2, a), as she would have limited her child's success and clashed with her tastes if s_2 would come true. Even if Alice feels a strong sense of responsibility for her child's success and happiness, she can't choose an action which will guarantee a compliance with an undefined profile preference. At least, she could try to avoid an outcome like (s_2, a), where she takes out a chances of success her child could have had without the intervention. In this case, a reasonable boundary to Alice's responsibility towards her child could be non-interference with a career her child could possibly choose or a taste she could display. Then, it is reasonable for Alice to focus on the worst outcome, because it's the one she wants to avoid.

In Section 4, the focus on the worst (and on the best) outcome will be formalised appealing to Hurwicz and Arrow's remarks on decision making under ignorance. Besides, the responsibility argument will justify the conservative and prudent feature of the proposed decision criterion.

IV. Justifying maximin

In this section I will put forward that, given the state of complete ignorance and the interpersonal feature of the decision context, maximin is a rational and convenient decision criterion for Alice's problem.

Following Milnor's axiomatic approach, I will argue that the column dupli-

cation axiom, which is not a generally reasonable requirement for a preference relation, holds in Alice's case. As shown by Arrow[2] and Hurwicz[3], column duplication implies that an agent's preferences over actions only depend on the worst and best consequences in any row of the matrix associated to the decision problem. This is easily verified if we identify an action with the row of consequences to which it leads. Next, following Binmore[4], it is verified that imposing convenient axioms over $\succeq_A$ implies that $argmax_{U_A} = max_{a_i}\{min_{s_k}(u_A(a,s))\}$ (Brafman and Tennenholtz 1997), where actions a_i are defined as the sets of consequences $c_i = (c_1, ..., c_k)$ attached to each action. Actually, the axioms justify a more general decision criterion, the so-called Hurwicz criterion. Nevertheless, I propose that the responsibility argument leads to a very conservative version of this criterion, which can be assimilated to maximin.

a. Milnor's axioms

I follow the formalisation proposed by Binmore[5] and Luce and Raiffa[6]:

- **Ordering** All actions must be completely ordered, i.e. $\succeq_A$ is a complete, symmetric and transitive relation over the sets of consequences attached to each action.
- **Symmetry** The ordering is independent of labeling of states and actions.
- **Strong domination** If each entry in the row representing the action $a_i$strictly exceeds the corresponding entry in the row representing the action a_j, then Alice strictly prefers a_i to a_j.
- **Continuity** Consider a sequence of decision problems with the same set of actions and states, in all of which $a_j \succ_A a_i$. If the sequence of matrices

[2] Arrow, 'Hurwicz's optimality criterion for decision making under ignorance'
[3] 'Hurwicz, Class of Criteria for Decision Making under Ignorance'
[4] Binmore, *Rational Decisions*
[5] Binmore, *Rational Decisions*
[6] Luce and Raiffa, *Games and Decisions*

of consequences converges, then its limiting value defines a new decision problem, in which $a_j \succeq a_i$.

- **Linearity** The order is not changed by linear utility transformations.
- **Row adjunction** Alice doesn't change her preferences between the old actions if a new action becomes available.
- **Column duplication** Alice doesn't change her preferences if a new column is appended that is identical to one of the old columns.

The column duplication axiom is not a generally justified property for decision making under uncertainty. As shown in Table 1, in the decision matrix on the right, which is obtained from the matrix on the left with column duplication, a_1 seems manifestly better than a_2, because it assures the payoff 1 in more states of nature. However, Hurwicz[7] argues that, if one is completely ignorant, then there is no reason to think that the true state is more likely to be one of the states s_2 to s_{10} than s_1. Anyway, this assumption has been criticized by Chernoff[8].

	s_1	s_2
a_1	0	1
a_2	1	0

	s_1	s_2	s_3	...	s_{10}
a_1	0	1	1	...	1
a_2	1	0	0	...	0

Table 1: Column duplication example

As Luce and Raiffa[9] notice, this axiom is used by Arrow and Hurwicz[10] to characterize the notion of complete ignorance in opposition to the principle of insufficient reason (or equiprobability). According to this principle, in the

[7] 'Hurwicz, Class of Criteria for Decision Making under Ignorance'
[8] Chernoff and Moses, 'Elementary Decision Theory'
[9] Luce and Raiffa, *Games and Decisions*
[10] Arrow and Hurwicz, 'An Optimality Criterion for Decision Making under Ignorance'

left matrix $p(s_1) = p(s_2) = 0.5$, while in the right matrix $p(s_1) = p(s_2) = ... = p(s_{10}) = 0.1$. But one can claim, as Arrow Hurwicz do, that adding a new state doesn't intuitively decrease the probability of the old ones. In their account, complete ignorance is to be interpreted as meaning that all states of nature have zero potential surprise. Then, dividing a state of nature into two would have no effect on the action chosen if the consequences of the action are the same under either substate. The sum of the probabilities attached to the states doesn't need to sum up to 1 as required by the principle of insufficient reason. In Alice's case, Arrow and Hurwicz's remarks seem sound. As we noticed in Section 2, Alice can't, and doesn't even need to, distinguish between all the possible states of nature. She only focuses of the outcomes that her actions will lead to in different relevant scenarios. So, adding a new column identical to an old one could be interpreted, in Alice's problem, as introducing a state where her child's attitudes and preferences doesn't affect the height enhancement in a different way. It can be regarded as a more detailed description of an already reckoned state, as well as a new one where the height enhancement would imply the same desiderable or undesiderable effects determined by an old state.

As a consequence of this property, Arrow[11] has proved the following result: if a criterion satisfies the axioms stated above, then it takes into account only the minumum and maximum utility associated with each action. This result conveniently characterize Alice's attitude towards the problem she is modeling. Given the interpersonal feature of the decision and the undefined prefence profile, she desires to act consistently with her child's preferences and to avoid to clash with them. In our example, she focuses on the best outcome (a, s_1) and on the worst one (a, s_2).

[11] Arrow, 'Hurwicz's optimality criterion for decision making under ignorance'

b. Hurwicz criterion and Maximin criterion

Axioms 1-6 and 8, taken together, imply the so called pessimism-optimism index criterion, or Hurwicz criterion.[12] This decision procedure prescribes the agent to look at a weighted combination of the best and worst outcome, following the insight of Arrow's result.

Hurwicz criterion[13] For actions $a_1, ..., a_n$, let m_i be the minimum and M_i the maximum of the utility numbers $u_{i_1}, ..., u_{i_n}$. Let a fixed number α between 0 and 1 be given. This index is called the pessimism-optimism index. To each a_i associate the index $\alpha m_i + (1-\alpha)M_i$, which is called the α-index of a_i. Of two acts, the one with higher α-index is preferred.

If $\alpha = 1$, Hurwicz criterion is equivalent to the maximin procedure:

Maximin criterion[14] To each action assign its security level as an index. Thus, the index for a_i is the minimum of the numbers $u_1, ..., u_n$. Choose that action associated whose associated index is maximum, i.e., choose the act that maximizes the minimum payoff.

The maximin criterion imposes a focus on the worst outcome, and forces the agent to avoid it. Manifestly, this implies not taking into account the best outcome which could be attained in the decision problem. That's why this criterion is labelled as ultra-conservative and pessimistic, and it is not regarded as a generally reasonable decision making procedure. As Chernoff[15] and Luce and Raiffa[16] point out, it would be paranoid and extremely risk-averse to adopt such a criterion, since it would mean to choose as if Nature were playing against us in a zero-sum game.

I suggest that in the particular context of Alice's case, this pessimistim can be justified appealing to the responsibility argument. In our example, it is

[12] Binmore, *Rational Decisions* 157-160.
[13] Luce and Raiffa, *Games and Decisions* 282-284
[14] Luce and Raiffa, *Games and Decisions* 278-280
[15] Chernoff and Moses, 'Elementary Decision Theory'
[16] Luce and Raiffa, *Games and Decisions*

reasonable to require Alice to avoid taking a decision which would limit her child's desire to undertake a career she could have perform otherwise (without the height enhancement). When making a decision on behalf of another individual, a prudent criterion and a maximun risk-aversion can be considered as istances of responsibility towards the person who can't express and uphold her preferences at the moment the choice is made.

c. The model

Let's call:

- $A_A = \{a, \neg a\}$ the set of the actions available to Alice: authorize or not authorize the height enhancement intervention. Actions a and $\neg a$ are defined as the subsets $c = \{c_{a1}, ..., c_{ak}\}$ and $c = \{c_{\neg a1}, ..., c_{\neg ak}\}$ of the consequences they bring about;

- $S = \{s_1, ..., s_k\}$ the set of relevant and distinguishable scenarios which Alice considers and which affects the outcomes of the decision;

- $C = \{c_{a1}, ..., c_{ak}\} \cup \{c_{\neg a1}, ..., c_{\neg ak}\}$ the set of the subsets of consequences attached to each action in A_A;

- $\succeq_A$ a preference relation defined over consequences in C, which satisfies Milnor's axioms 1-6 and 8.

Then, the optimal a_i is s. t. it maximizes $Min_{s_k}(U_A)$ (or s.t. it maximizes the α-index, where the pessmism-optimism index is 1).

Table 2 illustrates a possible model for Alice's problem, where:

- s_1 ='basketball player and finds above-average height good looking';

- s_2 ='jokey and finds above-average height good looking';

	s_1	s_2	s_3
a	5	2	0
$\neg a$	1	3	4

Table 2: Alice's decision problem

- $s_3 =$'jokey and finds above-average height awkward'

The maximin criterion prescribes Alice to select the maximum payoff over the minimum numbers in each row: $Max_A\{0,1\} = 1$. The adopted procedure imposes the agent not to authorise the height enhancement intervention, in order to avoid the worst possible outcome.

V. Precautionary Principle

In the previous sections we have justified the application of the maximin procedure within the decision model. The 'prudent' solution selected by Alice can be seen as an instance of the so-called Precautionary Principle (PP) as well. This general principle has emerged in the recent history of international environmental policy and law, in order to regulate the possibility of major human impacts on the global environment.[17] PP is a risk-managing procedure which imposes an attitude and consequently prescribes a convenient behaviour in such cases. Though many different versions of PP are put forward in specific contexts, I will refer in the following discussion to this standard statement:

> When an activity raises threats of harm to human health or the environment, precautionary measures should be taken even if some cause and effect relationships are not fully established scientifically. In this context the proponent of an activity, rather than the public, should bear

[17] See Myers, *Debating the Precautionary Principle, Science and Environmental Health Network*

> the burden of proof. The process of applying the precautionary principle must be open, informed and democratic and must include potentially affected parties. It must also involve an examination of the full range of alternatives, including no action.[18]

PP is variously criticised; for the sake of simplicity, I list here some of the main objections raised against it:[19] PP is vague and has conflicting definitions; if precaution applies to everything, it would stop all technology in its tracks; precaution calls for zero risk, which is impossible to achieve; precaution itself is risky: it will prevent us from adopting technologies that are actually safer; PP is anti-science. It goes beyond the aim of this paper to engage in answering these criticisms. I would simply point out that Alice's case is an intuitive, yet non trivial, example of application of PP. Moreover, the arguments supporting the application of maximin in Alice' problem (the uncertainty argument and the responsibility argument) are good points for PP as well. These features of the decision context could then represent suitable tools to delimit reasonable applications of PP, addressing the charges of vagueness and pessimism to the evaluation of the context itself.

The main idea behind PP is that frequently, in environmental or public health matters, a reliable procedure for risks quantification is not available. The decision context is far too complex for a well grounded balance between risks and benefits, since the mechanisms involved and the causal ties are not fully established. In such cases, the epistemic status can be described as Hurwicz's account of complete ignorance: risks exist as a mere possibility, even if we can't attach a probability to them. Environmental policy and public health worries also share the interpersonal character of Alice's case. The consequences of the choice will affect many other individuals, who can't take part personally in the decision making process (in the most Alice-like scenario, they belong to future

[18] Wingspread Statement on the Precautionary Principle, Jan. 1998

[19] Following Myers, *Debating the Precautionary Principle, Science and Environmental Health Network*

generations). Then, a focus on the worst outcome is normatively reasonable, and psychologically plausible, as this outcome is a risk we aren't willing to take, just like Alice doesn't want to take the risk of opposing her child's tastes. The instance of prudence required by PP is equivalent in this sense to maximin, as Gardiner[20] out too. Alice's case shows an example of justified application of PP (in its 'maximin version') subordinate to convenient features exhibited by the decision context. The state of complete ignorance and the strong responsibility towards other agents involved could set a reasonable ground for the normative instance of prudence advocated by PP.

References

[1] Arrow, K. J., 'Hurwicz's optimality criterion for decision making under ignorance', Technical Report 6, Stanford University Press, 1953.

[2] Arrow, K. J., and Hurwicz, L., 'An Optimality Criterion for Decision Making under Ignorance', in *Uncertainty and Expectations in Economics, Essays in Honour of G. L. S. Shackle*, Blackwell, 1972.

[3] Binmore, K., *Rational Decisions*, Section 9, Princeton University Press, 2009.

[4] Brafman, R. I. and Tennenholtz, M., 'On the Axiomatization of Qualitative Decision Criteria' in *American Association for Artificial Intelligence*, 1997.

[5] Chernoff, H. and Moses, L. E., *Elementary Decision Theory*, Dover Pub., 1954.

[6] Gardiner, S. M., 'A Core Precautionary Principle', in *The Journal of Political Philosophy*, Vol. 14, No. 1, 2006: 33-60.

[20] Gardiner, 'A Core Precautionary Principle'

[7] Hurwicz, L., 'A Class of Criteria for Decision Making under Ignorance', Cowles Commission Discussion Paper, Statistics No. 356, 1951.

[8] Luce, R. D. and Raiffa, H. *Games and Decisions*, J. Whiley, New York, 1957.

[9] Myers, N. *Debating the Precautionary Principle*, Science and Environmental Health Network, 2000.

Camilla Colombo is a MSc student in Philosophy and Public Policy at the London School of Economics and Political Science (2013-2014) and a MA student in Philosophy at the Scuola Normale Superiore di Pisa (2009-2014). She also holds a degree in Philosophy from University of Pisa. Her interests cover rational and collective decision-making, formal ethics and philosophy of biology. She hopes to continue research in the future. You can contact her at [camillafcolombo@gmail.com].