Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Multiple Beliefs, Dominance and Dynamic Consistency

Tommi Ekholm, Erin Baker

To cite this article:
Tommi Ekholm, Erin Baker (2021) Multiple Beliefs, Dominance and Dynamic Consistency. Management Science

Published online in Articles in Advance 15 Apr 2021

. https://doi.org/10.1287/mnsc.2020.3908

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–12
ISSN 0025-1909 (print), ISSN 1526-5501 (online)
http://pubsonline.informs.org/journal/mnsc

Multiple Beliefs, Dominance and Dynamic Consistency

Tommi Ekholm,[a,b] Erin Baker[c]

[a] Finnish Meteorological Institute, 00101 Helsinki, Finland; [b] Department of Mathematics and Systems Analysis, Aalto University, 02150 Espoo, Finland; [c] College of Engineering, University of Massachusetts, Amherst, Massachusetts 01003

Contact: tommi.ekholm@fmi.fi, https://orcid.org/0000-0002-9124-6289 (TE); edbaker@ecs.umass.edu, https://orcid.org/0000-0002-4139-9930 (EB)

Received: December 10, 2018
Revised: June 16, 2019; June 30, 2020
Accepted: September 17, 2020
Published Online in Articles in Advance: April 15, 2021

https://doi.org/10.1287/mnsc.2020.3908

Abstract. This paper investigates multiperiod decisions under multiple beliefs. We explore the dynamic consistency of both complete and incomplete orderings. We focus on a dominance concept that supports decision-making under multiple characterizations of uncertainty by ruling out strategies that are dominated across a set of beliefs. We uncover a distinction between two types of dynamic inconsistency, which we label fallacious and fallible inconsistency. Fallacious inconsistency occurs when an a priori optimal strategy is suboptimal in the second period, thus requiring the decision-maker to depart from the original strategy. Fallible inconsistency occurs when an a priori suboptimal second-period action ceases being suboptimal from the perspective of the second-period preferences. We introduce corresponding definitions of dynamic consistency and show that the two types of consistency are equivalent for complete orderings, but differ for incomplete orderings. Subjective expected utility is dynamically consistent and non-expected-utility decision rules, such as minmax, are not. We show that the dominance relation over beliefs falls between these two: it is immune to the more severe fallacious inconsistency, but not to the less problematic fallible inconsistency. We illustrate the method and concepts using a numerical example addressing a focal, real-world problem of risk and ambiguity regarding climate change.

History: Accepted by Ilia Tsetlin, decision analysis.
Funding: The work of T. Ekholm has been carried out with funding from the Academy of Finland [Decision number 311010]. The work of E. Baker has been supported by the University of Massachusetts Armstrong Professional Development Professorship.
Supplemental Material: The supplementary material is available at https://doi.org/10.1287/mnsc.2020.3908.

Keywords: uncertainty • ambiguity • multiple priors • dynamic decision making • dynamic consistency

1. Introduction

One key difficulty in decision-making is choosing between alternatives when their outcomes are uncertain.[1] If the outcomes' probabilities can be estimated reliably, the problem is a straightforward example of decision-making under risk. Decision-makers face a thornier problem if definite probabilities cannot be assigned for important uncertainties. We will refer to this as ambiguity. Such a situation arises, for example, if the decision-maker cannot aggregate multiple sources of information from multiple experts or statistical models (Cox 2012, Marinacci 2015). One way to approach this type of problem is to use multiple priors to account for the ambiguity (Gilboa and Schmeidler 1989). Each prior represents a subjective belief in a Bayesian sense, that is, a probability distribution.

A number of non-expected-utility decision rules have been proposed to address decision situations with ambiguity (Etner et al. 2012, Stoye 2012a, Gilboa and Marinacci 2016). We explicitly discuss a subset of these rules that are common in the literature. The maxmin expected utility criterion (Wald 1949, Gilboa and Schmeidler 1989) chooses an alternative relative to the least favorable distribution from a set of priors. Related is the α-maxmin expected utility (Ghirardato et al. 2004), which balances maxmin and maxmax rules. The minmax regret decision rule (Savage 1954) applies Wald's maxmin criterion to the possible regret faced by the decision-maker (Bell 1982, Loomes and Sugden 1982) and can be similarly extended for multiple beliefs.[2] Smooth ambiguity (Klibanoff et al. 2005) allows for the separation of preferences and beliefs by employing an ambiguity-aversion function in a parallel way as the risk-aversion function in expected utility. There is also a large body of literature on distributionally robust optimization applying sophisticated versions of these decision rules, primarily maxmin (e.g., Ben-Tal et al. 2013, Wiesemann et al. 2014, Bertsimas et al. 2019).

Another approach is to apply dominance rules, rather than decision rules. Dominance rules do not provide a complete ordering of alternatives, but rather rule out choices that are strictly worse than other available alternatives. Dominance rules have been used in real decisions, particularly in the context of multiple criteria (see examples in healthcare (Morton 2014), air traffic control (Grushka-Cockayne et al. 2008) and public housing (Johnson and Hurter 2000)). A number of authors have considered dominance relations over multiple beliefs, occasionally under the term admissibility: most prominently Bewley (2002), but also Wald (1949), Aumann (1962), Gilboa et al. (2010), Stoye (2012b), and Danan et al. (2016). Most of this literature aims at axiomatizing these rules and establishing preference relations. We focus on the framework introduced in Baker et al. (2020), who call the concept *belief dominance* and show that the nondominated set will include strategies resulting from subjective expected utility, maxmin expected utility, minmax regret, and smooth ambiguity.

Risk and ambiguity are particularly significant for decision problems dealing with long timeframes. In many practical cases, a long-term strategy consists of multiple consecutive decisions, rather than a one-shot decision. New information is often acquired between the decisions, enabling better-informed decisions in latter periods. A *strategy* refers to a complete specification of future actions at all foreseeable contingencies. We investigate such a multiperiod model, in which information arrives between the periods.

This intertemporal setting poses one additional challenge for decision-making: strategies can be dynamically inconsistent. Dynamic inconsistency occurs when a decision-maker chooses a first-period action with planned future actions following each possible outcome; but when the outcome is actually realized, they no longer want to follow the original plan. The problem is, if they knew they would not follow the plan in the future, they would have chosen something different in the first period (Strotz 1955). An inconsistent strategy lacks credibility, since if looked at carefully, it is apparent from the beginning that the plan will become suboptimal when re-evaluated in later periods. As planned actions are responses to foreseen observations, inconsistency is not due to either myopia or to the arrival of unexpected information, which can provide a rationale for changing plans. It is generally accepted that, for a decision model to have normative appeal, it needs to be dynamically consistent (Machina 1989). We note, however, that Siniscalchi (2009, 2011) has argued that dynamic consistency is simply one appealing criterion that can be traded off against other criteria such as ambiguity aversion.

It has been well established that non-expected-utility decision rules exhibit dynamic inconsistency, unless specific assumptions or formulations of the decision problem are made. Machina (1989) argues that dynamic inconsistency does not occur if the decision-maker acknowledges foregone opportunities in subsequent decisions. This solution, however, is problematic for prescriptive decision support, as it opens the possibility of accounting for sunk costs. A number of authors (Epstein and Schneider 2003, Wang 2003, Hayashi 2005, Maccheroni et al. 2006, Klibanoff et al. 2009) have established multiperiod ambiguity-averse preferences through a recursive formulation of utility. This forces the strategies to be dynamically consistent, but requires nontrivial assumptions to do so. For example, Epstein and Schneider (2003) and Hill (2020) restrict the permissible sets of beliefs. Al-Najjar and Weinstein (2009) argue that these recursive formulations assume away problems with dynamic consistency by ruling out partial resolution of ambiguity. Wang (2003) also notes that a nonrecursive formulation can be more convenient for some applications, with a lower computational burden, and that these two formulations yield the same result only under specific conditions.

While there exists abundant research on the dynamic consistency of the non-expected-utility rules discussed above, according to our knowledge only Ghirardato et al. (2008) have touched upon the dynamic consistency of dominance relations over beliefs. In this paper, we extend the static multiple-belief framework of Baker et al. (2020) to a multiperiod setting and explore the dynamic consistency of dominance rules in a two-period setting.

We point out here that there are two types of dynamic inconsistency, which we name fallacious and fallible. Fallacious inconsistency occurs when an a priori optimal strategy is suboptimal in the second period, thus requiring the decision-maker to depart from the original strategy. Fallible inconsistency occurs when an a priori suboptimal second-period action ceases being suboptimal from the perspective of the second-period preferences. In such case, the decision-maker *can* carry out the original strategy, but may also switch to a strategy that was originally suboptimal. A decision-maker would regret ex post committing to a fallacious strategy, but commitment removes the problem with a fallible strategy.

In the next section, we present a simple example illustrating some features of belief dominance and the two types of dynamic inconsistency in a two-period decision problem with multiple beliefs. Section 3 provides the formal definition of multiperiod belief dominance and defines two types of dynamic consistency. We show that the two types of consistency are equivalent for decision rules, which reflect complete orderings, but not for incomplete orderings, such as dominance rules. Last, we prove that belief dominance, for any set of priors and without a recursive formulation, avoids the more severe fallacious inconsistency, while nevertheless being subject to the less problematic fallible inconsistency. In Section 4, we apply the theory to a prime example of a real-world, long-term problem with ambiguity: cost-benefit analysis regarding climate change. We illustrate how belief dominance can narrow down to acceptable strategies under conflicting beliefs and the emergence of fallible dynamic inconsistency in this setting. In the last section, we conclude with a discussion of the merits and pitfalls of different approaches to dynamic decision-making under multiple beliefs.

2. An Illustrative Example

A virus outbreak has turned into a severe epidemic.[3] There are two strains of the virus, labeled A and B, that might have caused the epidemic. It will take one month to identify the strain. Health authorities have requested that several laboratories develop potential vaccines in two months, after which they will select one effective vaccine for a vaccination program.

A laboratory has three vaccines in development that could be effective against these two strains. Vaccine α is expected to be effective against the A strain, vaccine β against the B strain, and vaccine γ possibly effective against both strains. The laboratory can do research on only one vaccine at a time, but can change the focus of research after one month.

The laboratory's two leading virologists have estimated the probabilities of the strain being A or B and, for each vaccine, the probabilities that it is effective against each strain, conditional on the time put into its research.[4]

The virologists have differing beliefs on the likelihood of the strain and the effectiveness of the vaccines, shown by the virologists' expressed probabilities presented in Table 1. We assume that a vaccine is either effective or ineffective. The effectiveness of each different vaccine is assumed to be stochastically independent and not mutually exclusive.

Based on this information, the laboratory's manager is contemplating a research strategy for the next two months. The manager wishes to maximize the expected number of effective vaccines at the end of this period, taking into account the ambiguity expressed in the two virologists' beliefs.

A *strategy* is a triplet $x = (x_1, x_{2,A}, x_{2,B})$, with each element $x_{i\in\{1,(2,A),(2,B)\}} \in \{\alpha, \beta, \gamma\}$ reflecting a decision to research a given vaccine for one month. The triplet's elements correspond to the first-period research focus (x_1) and the second-period focus, dependent on whether strain A or B is observed ($x_{2,\mathrm{A}}$ and $x_{2,\mathrm{B}}$). For example, strategy (β, α, β) starts with research on vaccine β, and proceeds with research on vaccine α if the strain turns out to be A, or on vaccine β if the strain turns out to be B. There are $3^3 = 27$ different strategies in total.

The expected values for the number of effective vaccines are presented in Figure 1 for all strategies and both beliefs. See Appendix A for a formalization and details regarding the calculations. We discuss the optimal strategies under multiple-decision rules, including subjective expected utility (SEU), smooth ambiguity, maxmin, and minmax regret over the two beliefs.

Seven strategies, marked with yellow in Figure 1, are nondominated across the two beliefs in the first period: there is no strategy that has a higher expected number of vaccines under both beliefs. The set of nondominated solutions includes the maxmin, minmax regret, SEU, and smooth ambiguity solutions (Baker et al. 2020). The optima for the maxmin and minmax regret, (γ, α, γ) and (γ, α, β), respectively, are highlighted in Figure 1. The SEU and smooth ambiguity optima depend on the weighting or subjective probability over the two virologists and can be any of the nondominated strategies except (γ, α, β).

Table 1. The Probabilities Expressed by Each Virologist of Whether the Strain is A and of the Vaccines' Effectiveness Against Each Strain, Conditional on the Research Effort Put Toward that Vaccine

		Research effort	Virologist 1	Virologist 2
Probability of strain being A?			90%	10%
Probability of vaccine being effective against strain A		α	60%	40%
	with one month of research?	β	20%	10%
		γ	10%	20%
		(α, α)	85%	80%
	with two months of research?	(β, β)	40%	55%
		(γ, γ)	30%	80%
Probability of vaccine being effective against strain B		α	10%	0%
	with one month of research?	β	30%	50%
		γ	40%	20%
		(α, α)	40%	10%
	with two months of research?	(β, β)	80%	65%
		(γ, γ)	65%	80%

Note. One letter stands for one month of research on the stated vaccine, and two letters stand for two months.

Figure 1. (Color online) The Expected Number of Effective Vaccines Against the Unknown Virus Strain Under the Two Virologists' Beliefs

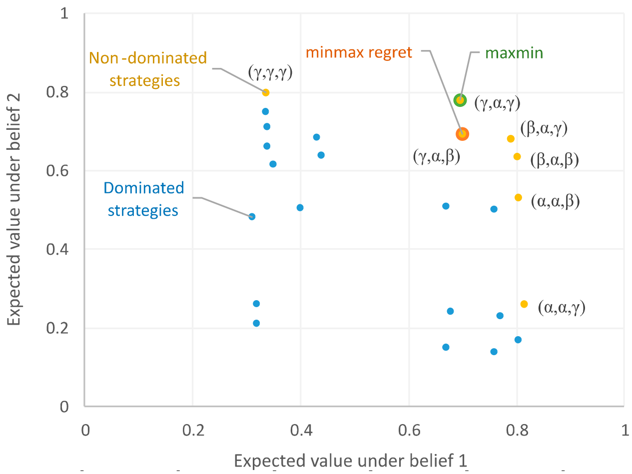

Notes. The labels show the values of the triplet $(x_1, x_{2,A}, x_{2,B})$ for nondominated points.

The nondominated set, as a whole, goes beyond any of the individual solutions, providing a comprehensive perspective on the trade-offs the manager might make between the two beliefs.

Consider that the manager uses either maxmin or minmax regret as the decision rule. Both suggest researching vaccine γ during the first month. If the strain turns out to be B, then according to the maxmin strategy (γ, α, γ), further research should be done on γ. However, if the manager recalculates the maxmin strategy after learning the strain is B, the optimal action is to research vaccine β. Research on γ would be suboptimal, although this was suggested as optimal after observing the strain being B by the original strategy. Similar inconsistency takes place with the minimax regret decision rule, which prescribes strategy (γ, α, β) in the first period. After finding out that the strain is B, the recalculated minimax regret action is to research vaccine γ. It would now be suboptimal to research vaccine β, which was implied by the original minimax regret strategy.

If the manager follows the same decision rule in both periods, the second-period optimum would contradict the optimality of the first-period strategy. The decision rule reverses preferences between the alternatives when evaluated in different periods and is therefore inconsistent over time.

Belief dominance does not provide a decision rule, but can be used to narrow down to a set of acceptable, nondominated alternatives. This selection behaves in a more consistent—though not perfectly consistent—manner than the two decision rules portrayed above.

For every nondominated strategy, the associated second-period actions continue to be nondominated if recalculated in the second period. The manager can thus continue to follow the original strategy. However, in some cases the manager might find that a second-period action that was initially dominated has become nondominated. For example, if following the nondominated strategy (β, α, β) and finding out that the strain is A, the second-period action of β becomes nondominated, although all strategies of the form $(\beta, \beta, x_{2,B})$ were dominated in the first period. We will show in the following section that these are inherent features of non-belief-dominated strategies.

In some cases, applying belief dominance may reveal a good solution without forcing stakeholders to agree on beliefs. While this is not the case in this example, it does lead to considerable narrowing down of alternatives: from 27 to 7. It also provides a method for visualizing the tradeoffs. For example, Figure 1 reveals a small loss under Belief 1 when moving from (α, α, γ) to (β, α, γ), but under Belief 2 the associated improvement is very large. Even a decision-maker with high confidence in Virologist 1 may be willing to compromise. Further, the approach can illustrate why disagreeing experts suggest conflicting strategies and help build acceptance for a middle-ground solution.

3. Definitions and Properties

Decisions are made in two periods: 1 and 2. The decision-maker considers possible outcomes that result from these decisions and a random variable Z with possible realizations $z \in \mathcal{Z}$. Between the periods, the decision-maker observes a signal Y, a random variable that is related to Z with possible realizations $y \in \mathcal{Y}$. The random variables are defined on a probability space (Ω, A, f), where $\Omega = \mathcal{Y} \times \mathcal{Z}$ is the sample space, A is the σ-algebra of events, and f is the probability measure. An outcome is a realization of both random variables $(y, z) \in \Omega$.

To present the decision problem in extensive form, strategy $x = (x_1, x_2(Y)) \in X$ comprises first-period action $x_1 \in X_1$ and second-period actions $x_2(y) \in X_{2,x_1,y}$ following each possible realization $Y = y$, with $y \in \mathcal{Y}$. It is assumed that the second-period decision space $X_{2,x_1,y}$ can depend on the first-period action x_1 and the observation of Y. The random variables are assumed to be exogenous, that is, independent of the chosen strategy x.

The decision-maker wishes to maximize expected utility, but is faced with conflicting subjective beliefs on the probabilities of Y and Z that she deems plausible. The set of beliefs the decision-maker considers is denoted with Φ. The conditional probabilities of Z after $Y = y$ has been observed, $f_Z(z|y)$, can be determined through Bayesian updating through the marginal distributions for the realizations y and z: $f(z, y) = f_Z(z|y) f_Y(y)$. In the case where $f_Y(y) = 0$ for some particular $y \in \mathcal{Y}$, we define $f(z, y) = 0$ and say that

$f_Z(z|y)$ is undefined. Define the set Φ_y as the subset of prior beliefs f in Φ for which $f_Z(z|y)$ are well defined.

The decision-maker's preferences are represented with a utility function $U(x_1, x_2, z)$. In this case, U is a general objective function that can include risk preferences, calculations tying alternatives to outcomes, and aggregation of multiple criteria into a univariate measure of utility. Given a belief f, the first-period expected utility for strategy x is represented by V:

$$V(x,f) = \int_{y \in \mathcal{Y}} \int_{z \in \mathcal{Z}} U(x_1, x_2(y), z) f(z, y)\, dz dy. \tag{1}$$

In the second period, following the first-period decision x_1 and observation of $y \in \mathcal{Y}$, the expected utility for a belief $f \in \Phi_y$ is

$$V_2(x_1, x_2, y, f) = \int_{z \in \mathcal{Z}} U(x_1, x_2, z) f_Z(z|y) dz. \tag{2}$$

3.1. Definition of Belief Dominance

In the first period, strategy $x^* = (x_1^*, x_2^*(Y))$ belief-dominates strategy $x = (x_1, x_2(Y))$ if

$$V(x^*, f) \geq V(x, f) \;\; \forall f \in \Phi, \tag{3}$$

with the inequality being strict for some $f \in \Phi$. This represents the preference $x^* \succ x$. The decision-maker is indifferent between two strategies, written $x^* \sim x$, if their expected utilities are equal $\forall f \in \Phi$. The relation $\succcurlyeq$ implies that either of these holds, and this relation is reflexive and transitive—hence a preorder. Belief dominance, however, is not a complete order, as some strategies are not comparable.[5] We write $x^* \not\prec x$ to denote that x^* is not strictly dominated by x.

In the second period, after x_1 has been carried out and y observed, action x_2^* belief-dominates action x_2 if

$$V_2(x_1, x_2^*, y, f) \geq V_2(x_1, x_2, y, f) \; \forall f \in \Phi_y, \tag{4}$$

with the inequality being strict for some $f \in \Phi_y$. We write this as $x_2^* \succ_{x_1, y} x_2$.

The purpose of belief dominance in decision support is to narrow down the choices to strategies that are *nondominated*: strategy x^* is nondominated if there does not exist a strategy $x \in X$ that dominates x^*. A dominated strategy is not a good choice, because another strategy exists that is at least as good under all considered beliefs and strictly better under at least one belief. Comparisons among nondominated strategies, on the other hand, involve trade-offs between expected outcomes predicted by different beliefs.

All nondominated strategies for a given problem and set of beliefs Φ form the *nondominated set*. The set of nondominated strategies contains the optimal solution to any convex combination of the beliefs, but may contain other strategies as well (Baker et al. 2020).

3.2. Two Types of Dynamic Inconsistency

Let us now formally define the two types of inconsistency that were encountered in the example above. Let $\preccurlyeq$ be a preference relation between possible strategies. Throughout this section, we assume that $\preccurlyeq$ is a preorder (not necessarily complete). Let X^* represent the set of a priori "optimal" strategies, defining "optimal" to mean that a strategy is not strictly preferred by any other available strategy in the first period: $x^* \in X^* \Leftrightarrow x^* \not\prec x, \; \forall x \in X$. Thus, with a complete ordering, x^* would be optimal in the more usual sense; with an incomplete ordering, x^* would be nondominated.

Assume that a first-period action x_1 from a strategy $x = (x_1, x_2(Y))$ is carried out and a signal y is observed. Let $x_2(Y \neq y)$ denote the strategy's set of second-period actions that would follow other possible signals $Y \neq y$. We then write strategy x in the form $(x_1, x_2(y), x_2(Y \neq y))$. Let $\preccurlyeq_{x_1, y}$ be the second-period preference ordering, following a first-period action x_1 and observed signal y. Let $X^*_{2, x_1, y}$ be the corresponding set of optimal second-period actions: $x_2^* \in X^*_{2, x_1, y} \Leftrightarrow x_2^* \not\prec_{x_1, y} x_2, \; \forall x_2 \in X_{2, x_1, y}$.

In this setting, we define two types of inconsistency:

1. Fallacious inconsistency: A decision problem exhibits fallacious inconsistency if a second-period action that is a part of a strategy in the first-period optimal set is not in the second-period optimal set. Formally, $\exists x^* \in X^*$ with $x_2^*(y) \notin X^*_{2, x_1^*, y}$ for some $y \in \mathcal{Y}$.

2. Fallible inconsistency: A decision problem exhibits fallible inconsistency if a second-period optimal set contains an action that was not in any strategy of the first-period optimal set, despite the first-period action being optimal. Formally, $\exists x_2 \in X^*_{2, x_1^*, y}$ for some $y \in \mathcal{Y}$ such that $\forall \hat{x}_2(Y \neq y) \in \prod_{\hat{y} \neq y} X_{2, x_1^*, \hat{y}}$ $\exists x \in X$ so that $x \succ (x_1^*, x_2, \hat{x}_2(Y \neq y))$.

That is, with fallacious inconsistency a decision-maker is *required* to depart from her original strategy, because carrying out that strategy is suboptimal in the second period. This is the problem most often illustrated with decision rules exhibiting dynamic inconsistency. For example, Machina (1989) presents dynamic inconsistency as fallacious inconsistency, while Epstein and Schneider (2003) provide an axiomatization of this type of consistency for complete preference orderings.

By contrast, we are not aware of literature that has expressly discussed fallible inconsistency. A decision-maker *can* carry out the strategy devised in the first period with fallible inconsistency, as the strategy remains in the optimal set in the second period. But she can also find second-period actions that seem optimal at that point, although they were not optimal in the first period. A preference ordering that can exhibit fallible inconsistency does not force the decision-maker off-course, but can be seen as problematic if the decision-maker willingly departs from her original strategy.

Fallible inconsistency occurs only when (1) a strategy x is dominated solely by strategies $\hat{x}$ that have a differing first-period action $\hat{x}_1 \neq x_1$, (2) there are non-dominated strategies that contain the first-period action x_1, and (3) there exists a branch $Y = y$ where the second-period action $x_2(y)$ is not dominated by the available actions $\tilde{x}_2 \in X_{2,x_1,y}$. If action x_1 is taken in the first period, then the strategy $\hat{x}$ is no longer available in the second period, and therefore, strategy x ceases to be dominated in the $Y = y$ branch.

The difference between fallacious and fallible inconsistency can be framed through commitment.[6] A decision-maker who acknowledges fallacious inconsistency would always be willing to pay, in the first period, some small amount for commitment (see, e.g., Strotz 1955, Thaler and Shefrin 1981, Gul and Pesendorfer 2001, Battigalli et al. 2019). Under fallible inconsistency, in contrast, no incentive is needed to follow the original strategy, and the decision-maker would thus not be willing to pay for commitment.

3.3. Dynamic Consistency of Preference Orderings

Dynamic inconsistency is a phenomenon that can, but is not bound to, happen with certain preference orderings in specific decision-making settings. To characterize preference orderings generally, we introduce associated definitions of consistency:

1. Weak consistency: If $x, \tilde{x} \in X$, $x_1 = \tilde{x}_1$, and $x_2(Y \neq y) = \tilde{x}_2(Y \neq y)$ for some $y \in \mathcal{Y}$, then
   i. $x_2(y) \prec_{x_1,y} \tilde{x}_2(y)$ implies $x \prec \tilde{x}$, and
   ii. $x_2(y) \preccurlyeq_{x_1,y} \tilde{x}_2(y)$ implies $x \preccurlyeq \tilde{x}$.
2. Strong consistency: Given $x, \tilde{x} \in X$, $y \in \mathcal{Y}$, and $\hat{x}_2 \in X_{\tilde{x}_1,y}$,
   a. if $\tilde{x} \not\prec x$ and $\hat{x}_2 \not\prec_{\tilde{x}_1,y} \tilde{x}_2(y)$, then $(\tilde{x}_1, \hat{x}_2, \tilde{x}_2(Y \neq y)) \not\prec x$;
   b. if $\tilde{x} \succcurlyeq x$ and $\hat{x}_2 \succ_{\tilde{x}_1,y} \tilde{x}_2(y)$, then $(\tilde{x}_1, \hat{x}_2, \tilde{x}_2(Y \neq y)) \succ x$; and
   c. if $\tilde{x} \sim x$ and $\hat{x}_2 \sim_{\tilde{x}_1,y} \tilde{x}_2(y)$, then $(\tilde{x}_1, \hat{x}_2, \tilde{x}_2(Y \neq y)) \sim x$.

The definition for dynamic consistency often found in literature (e.g., Epstein and Schneider 2003) corresponds to our definition of weak consistency. As the names suggest, strong consistency implies that weak consistency also holds. The weak and strong consistency, respectively, rule out fallacious and fallible inconsistency. We prove below that belief dominance conforms to weak consistency and thus can only exhibit fallible inconsistency.

The distinction between these two has not been pointed out previously, we think, because they are equivalent in the case of decision rules with complete preference orderings, as we show below. Strong consistency additionally covers relations between incomparable alternatives, which is relevant for dominance rules with incomplete preference ordering, and provides implications for relations between strategies having differing first-period actions, which is necessary to rule out fallible inconsistency.

Theorem 1. *Strong consistency implies weak consistency.*

Proof. Assume that strong consistency holds but weak consistency does not. Consider that the weak consistency assumptions $x_1 = \tilde{x}_1$ and $x_2(Y \neq y) = \tilde{x}_2(Y \neq y)$ hold. Let us analyze separately the two implications of weak consistency following the action x_1 and observation y. Consider first a violation of (i) and then, as a remaining case, a violation of (ii) if the second-period actions are equally preferable, as strict preference was considered already under (i).

Case (i): $\tilde{x}_2(y) \prec_{\tilde{x}_1,y} x_2(y)$.
Assume that weak consistency is violated: although $\tilde{x}_2(y) \prec_{\tilde{x}_1,y} x_2(y)$, it holds that $\tilde{x} \not\prec x = (\tilde{x}_1, x_2(y), \tilde{x}_2(Y \neq y))$. But strong consistency condition (b) says that if $\tilde{x} \sim \tilde{x}$ and $x_2(y) \succ_{\tilde{x}_1,y} \tilde{x}_2(y)$, then $(\tilde{x}_1, x_2(y), \tilde{x}_2(Y \neq y)) = x \succ \tilde{x}$. This contradicts the violation of weak consistency with $\tilde{x} \not\prec x$.

Case (ii): $\tilde{x}_2(y) \sim_{\tilde{x}_1,y} x_2(y)$.
The proof is similar and can be found in the Appendix. Therefore, strong consistency implies weak consistency. ■

For a complete preordering, weak also implies strong.

Theorem 2. *For complete preorders $\succcurlyeq$ the two definitions of dynamic consistency are equivalent.*

Proof. We show here that weak consistency implies strong consistency for complete orderings, for which $\tilde{x} \not\prec x \Leftrightarrow \tilde{x} \succcurlyeq x$.

Assume that implication (a) of strong consistency does not hold. Then the following will be true:

a. $\tilde{x} \succcurlyeq x$,
b. $\hat{x}_2 \succcurlyeq_{\tilde{x}_1,y} \tilde{x}_2(y)$,
c. $\hat{x} = (\tilde{x}_1, \hat{x}_2, \tilde{x}_2(Y \neq y)) \prec x$.

If weak consistency holds, (b) above implies $\hat{x} \succcurlyeq \tilde{x}$. Combining this with (c) gives $\tilde{x} \preccurlyeq \hat{x} \prec x$, which is a contradiction with (a). Similar reasoning holds for the two latter implications of strong consistency, and therefore, weak consistency implies strong consistency for complete and transitive orderings. Combining this with Theorem 1 completes the proof. ■

Here we state two theorems showing the relationships between the two definitions of consistency and the two definitions of inconsistency. The proofs are provided in the appendix.

Theorem 3. *A preference ordering $\preccurlyeq$ that has the property of weak consistency cannot exhibit fallacious inconsistency.*

Proof. See the appendix.

Theorem 4. *A preference ordering $\preccurlyeq$ that has the property of strong consistency cannot exhibit fallible inconsistency.*

Proof. See the appendix.

3.4. Belief Dominance Is Weakly Consistent

Belief dominance has the property of weak consistency, thus avoiding fallacious inconsistency. The proof for this is provided for discrete random variables, so that each possible signal $y \in \mathcal{Y}$ has a positive probability mass.[7] We note that even expected utility can exhibit fallacious inconsistency for realized signals $\hat{y}$ that have zero subjective probability.

Theorem 5. Belief dominance is weakly consistent.

Proof. Consider $x = (x_1, x_2(Y))$ and $\tilde{x} = (x_1, \tilde{x}_2(y), x_2(Y \neq y))$. Assume that x_1 is carried out, that a signal y is observed after period one, and that the second-period decision $\tilde{x}_2(y)$ belief-dominates $x_2(y)$. This means that

$$\sum_{z \in \mathcal{Z}} U(x_1, x_2(y), z) f_Z(z|y) \leq \sum_{z \in \mathcal{Z}} U(x_1, \tilde{x}_2(y), z) f_Z(z|y) \;\; \forall f \in \Phi_y,$$

with the inequality being strict for at least one $f \in \Phi_y$.

Since $f_Y(y) \neq 0 \; \forall f \in \Phi_y$, we can multiply both sides by $f_Y(y)$, which gives

$$\sum_{z \in \mathcal{Z}} U(x_1, x_2(y), z) f(z, y) \leq \sum_{z \in \mathcal{Z}} U(x_1, \tilde{x}_2(y), z) f(z, y) \;\; \forall f \in \Phi_y.$$

Since $f(z, y) = 0$ for any $f \in \Phi$ that has $f_Y(y) = 0$, the above implies

$$\begin{aligned} \sum_{z} U(x_1, x_2(y), z) f(z, y) + \sum_{(z, \hat{y} \neq y) \in \Omega} U(x_1, \tilde{x}_2(\hat{y}), z) f(z, \hat{y}) \\ \leq \sum_{z} U(x_1, \tilde{x}_2(y), z) f(z, y) \\ + \sum_{(z, \hat{y} \neq y) \in \Omega} U(x_1, \tilde{x}_2(\hat{y}), z) f(z, \hat{y}) \;\; \forall f \in \Phi, \end{aligned}$$

with the inequality being strict for at least one $f \in \Phi$. This is equal to the definition of $\tilde{x}$ belief-dominating x, indicating that the implication (i) of weak consistency holds. A similar argument applies for case $\tilde{x}_2(y) \sim_{x_1, y} x_2(y)$, and when combined, these cover the implication (ii) of weak consistency. Belief dominance is therefore weakly consistent. ■

Thus, belief dominance is not fully dynamically consistent, but it does avoid the most egregious inconsistency. As the solutions to many non-expected-utility decision rules are contained in the nondominated set (Baker et al. 2020), it may seem surprising that these rules are subject to fallacious inconsistency when belief dominance is not. What this means is that a second-period action of a non-expected-utility maximization might be suboptimal in terms of that specific decision rule, but the action is still nondominated in the second period.

4. A Numerical Example: Cost-Benefit Analysis on Climate Change

Climate change is a large-scale problem requiring long-term strategies and involving ambiguity. Individuals and entities with conflicting beliefs have long argued over how much should be done to mitigate climate change. The combination of ambiguity and long-term dynamics make this problem particularly difficult and controversial.

We focus here on climate change mitigation: the reduction of greenhouse gasses compared with a business-as-usual situation. Reducing emissions incurs costs, but also lessens the impacts that a changing climate has on society and ecosystems. Ideally, mitigation strategies would minimize the sum of mitigation costs and climate damages. Pioneering analyses in this cost-benefit setting were done by Nordhaus (1992, 2017) with his DICE model.

However, two critical challenges for crafting an optimal mitigation strategy are the ambiguity around how much the climate will warm under various mitigation strategies (Millner et al. 2013) and the severity of damages from a given level of warming (Diaz and Moore 2017). A number of papers have considered these uncertainties in a probabilistic manner (see Lemoine and Rudik (2017) for a review). We build on a sensitivity analysis on the combination of these two sources of ambiguity by Ekholm (2018), using a simplified version of his dynamic cost-benefit model. Yet, we go beyond the sensitivity analysis of probabilities, identify dynamic strategies that are nondominated under diverse beliefs, and illustrate the fallible inconsistency that arises with the nondominated strategies.

4.1. Model Description

A *strategy* here is the combination of a near-term emissions target and a long-term emissions target conditioned on information received after the near-term action. Specifically, the near-term action x_1 represents a global emission target for year 2030, such as those considered currently under the United Nations' Paris Agreement. The long-term action x_2 corresponds to an emission target for 2070, representing mitigation actions during the latter half of the century. We assume climate damages are realized around 2100. A more detailed description of the model is provided in the supplementary material.

In the absence of any action, a baseline level of emissions would take place, and reducing emissions from this level incurs costs. We represent the cost of reducing emissions below this baseline level as a convex function of the emission reduction level. We adopt the specific cost function from Ekholm (2018). The costs in 2070 reflect future technological development; thus, they are lower than the costs in 2030. The economic value of climate damages is modelled as a power function of temperature change, as in the widely used DICE model (Nordhaus, 1992, 2017). The objective is to minimize the expected, discounted sum of mitigation costs and climate damages.

We consider ambiguity in two dimensions: climate sensitivity and damages. Climate sensitivity (Cs) is a measure of how much the temperature will increase due to a doubling of atmospheric CO_2 concentration. There is significant disagreement about the probability distribution over this value. We define two beliefs that correspond to discretized versions of the symmetric and lognormal distributions presented by Knutti and Hegerl (2008).

For climate damages, we incorporate ambiguity in the damage function's exponent, the motivation being that the quadratic representation in DICE was deemed highly uncertain early on (Nordhaus 1992) and has received notable criticism (e.g., Pindyck 2013). We adopt two extreme beliefs around the exponent, and note that they are speculative and used for illustrative purposes. We label the two beliefs as "unconcerned" and "alarmist," with the former believing the highest damages are unlikely, while the latter has a converse view. The probabilities for climate sensitivity and damages are assumed to be independent. The combination of two beliefs for both parameters leads to four beliefs in total.

The two-dimensional random variable Z comprises the values of climate sensitivity and the damage exponent. Each dimension of the random variable has three possible values, leading to nine possible outcomes. A signal Y is received between decisions x_1 in 2030 and x_2 in 2070. This signal reflects information from new measurements, modelling, and impact analyses and narrows down the possible values of Z by ruling out either the highest or the lowest value for each of the two dimensions independently, with equal probability, regardless of the underlying distribution for Z. If the highest parameter value is excluded, we denote this as "low-to-medium," and we denote it "medium-to-high" if the lowest parameter value is excluded. Four different signals are therefore possible. This structure is illustrated in Figure 2. The numerical values for these branching probabilities are presented in Table 2.

Given a particular belief, the resulting optimization problem is convex, with both a convex decision space and a convex and continuous objective function under all of the considered parameters. Consequently,

Figure 2. (Color online) Decision Tree for the Climatic Cost-Benefit Problem

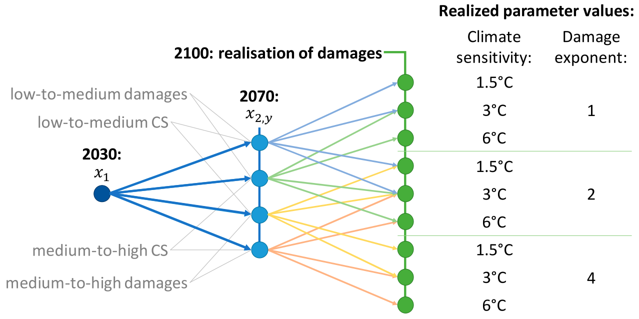

Notes. Decisions in 2030 are made before learning. Between 2030 and 2070 one out of four possible signals is received, indicating which values of Cs and damage exponent are eliminated. The shapes in 2070 denote the information states and correspond to those used later in Figure 3. The final outcomes of climate damages are realized in 2100, with three possible values for Cs and the damage exponent.

all nondominated strategies are solutions to an optimization problem that employs some convex combination of the four beliefs (Geoffrion 1968). The non-belief-dominated set can thus be approximated by using a grid of weightings w_i on the four-dimensional simplex over the four beliefs $\{(w_1,w_2,w_3,w_4)\in[0,1]^4|w_1+w_2+w_3+w_4=1\}$. The grid's intervals were set at 5%, leading to 1,771 different weightings between the beliefs.

4.2. Nondominated Mitigation Strategies and Dynamic Consistency

Figure 3 presents the nondominated set as the correspondence between 2030 and 2070 emission levels, separately for each second-period branch (see Figure S2 in the supplementary material for a visualization of the strategies over time). The optimal strategies for each of the four "pure" beliefs are denoted with markers. In 2030, the belief with a symmetric distribution for climate sensitivity and unconcerned view on damages leads to the highest optimal emissions, while the skewed distribution for climate sensitivity combined with the alarmist stance on damages leads to the lowest. The beliefs on damages have far greater impact on the optimum than beliefs on climate sensitivity.

The dark blue area illustrates the set of strategies that are nondominated from the perspective of the first period. The light blue area represents the range of second-period actions that are nondominated in the second period, but not in the first. The first-period nondominated set is a subset of the second-period nondominated set, showing that fallacious inconsistency is avoided, consistent with Theorem 5. That is, any nondominated strategy chosen in the first period is still nondominated in the second period. However, for each first-period action there are second-period actions that are nondominated in the second period,

Table 2. Beliefs Used in the Climatic Cost-Benefit Analysis

	Symmetric	Skewed
P(Y_{CS} = low-to-medium *Cs*)	50%	50%
P(Y_{CS} = medium-to-high *Cs*)	50%	50%
P(Z_{CS} = 1.5°C | Y_{CS} = low-to-medium *Cs*)	21%	47%
P(Z_{CS} = 3°C | Y_{CS} = low-to-medium *Cs*)	79%	53%
P(Z_{CS} = 3°C | Y_{CS} = medium-to-high *Cs*)	77%	63%
P(Z_{CS} = 6°C | Y_{CS} = medium-to-high *Cs*)	23%	37%
	Unconcerned	Alarmist
P(Y_D = low-to-medium damages)	80%	20%
P(Y_D = medium-to-high damages)	20%	80%
P(Z_D = 1 | Y_D = low-to-medium damages)	80%	20%
P(Z_D = 2 | Y_D = low-to-medium damages)	20%	80%
P(Z_D = 2 | Y_D = medium-to-high damages)	80%	20%
P(Z_D = 4 | Y_D = medium-to-high damages)	20%	80%

Notes. With the notation from Section 2, Y_{CS} and Y_D denote the signals on climate sensitivity and damages, and Z_{CS} and Z_D denote the random variables of the parameter values. The values shown are the probabilities of Y_{CS}, Y_D, Z_{CS}, and Z_D associated with each belief.

Figure 3. (Color online) Nondominated Emissions Levels in the First Period (*x*-Axis) and in the Second Period (*y*-Axis) in Different Signals Regarding Climate Sensitivity (Columns) and Damages (Rows)

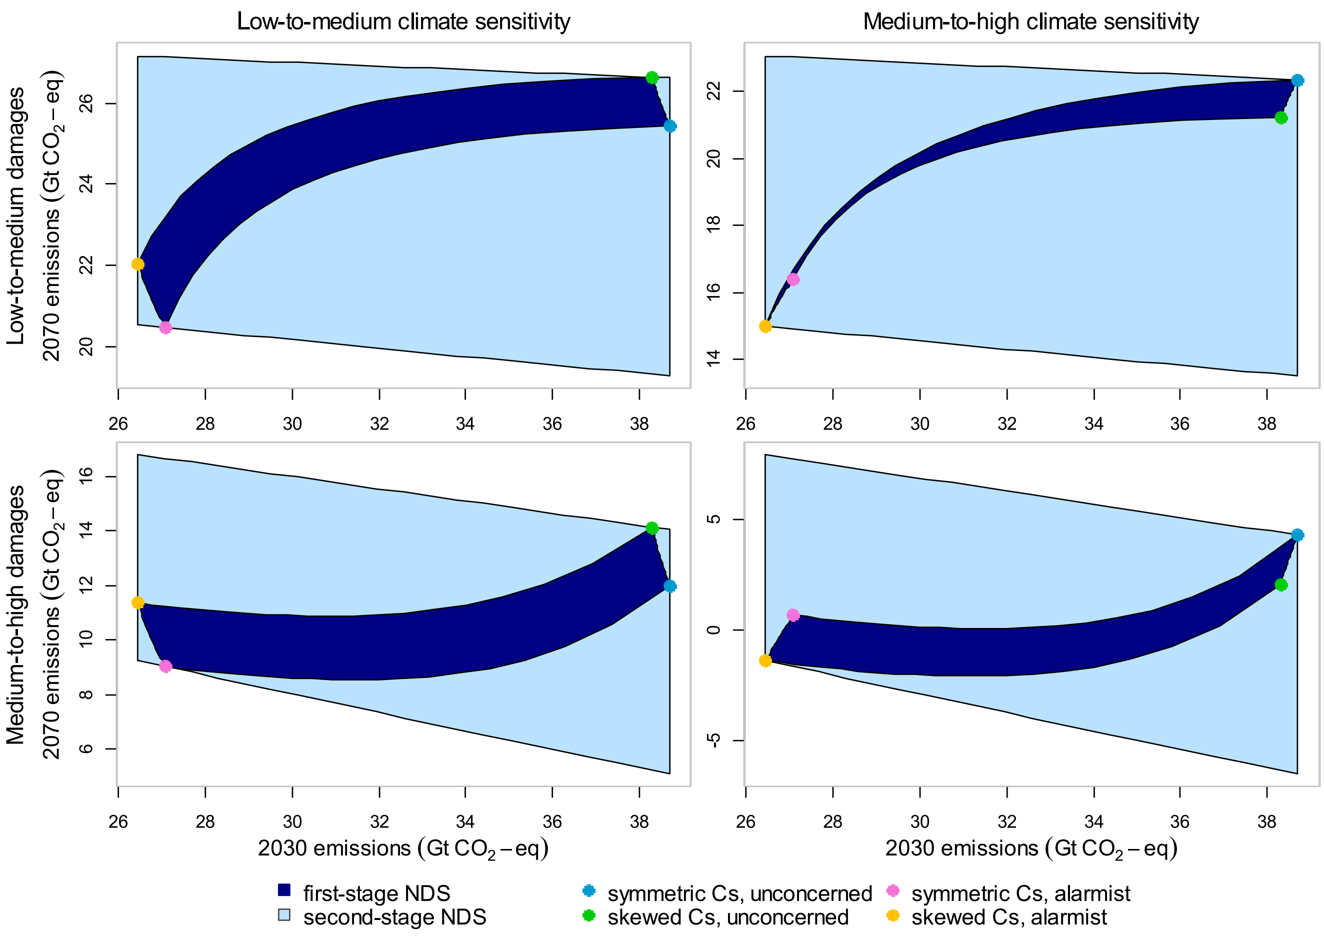

Notes. The dark area represents strategies in the first-period nondominated set (NDS). The light area represents for each first-period action the second-period actions that are nondominated in the second period.

but were dominated in the first period, implying that fallible inconsistency is present.

As an example, consider a case where the first-period decision x_1 is to limit emissions to 32 Gt. The nondominated strategies that contain this first-period action have second-period emissions ranging from roughly 9 Gt to 11 Gt if the signal Y indicates medium-to-high damages and low-to-medium Cs (the lower left-hand chart in Figure 3). However, once in that branch in the second period, emissions as high as 15 Gt are now nondominated. This is despite that, from the first-period perspective, all combinations of 32 Gt in the first period and 15 Gt in that particular second-period branch were strictly dominated, specifically, by strategies that have first-period emissions higher than 32 Gt and second-period emissions lower than 15 Gt in that branch. But since such dominating strategies are no longer available due to the chosen first-period action of 32 Gt, the second-period nondominated set has expanded. Similarly, the second-period nondominated set reaches down to 8 Gt, because by choosing 32 Gt in the first period, we have forgone dominating strategies that involve lower first-period emissions than the chosen 32 Gt.

More generally, fallible inconsistency introduces nondominated second-period actions that work against the chosen first-period action. First-period actions in the low-end of the nondominated set, for example, are more representative of the “alarmist” belief. The fallible inconsistency in this case introduces higher emissions to the second-period nondominated set, more representative of the “unconcerned” case. The opposite holds for strategies with high first-period emissions.

Note that we do not employ Bayesian updating between the beliefs, as dominance is evaluated against the whole set of beliefs in both periods. Fallible inconsistency does not arise due to a lack of updating, but due to the unavailability of some foregone strategies as a result of the chosen first-period action.

5. Conclusions and Discussion

Our contribution has been on two fronts: practical and theoretical. First, we have expanded the belief dominance concept to multiperiod problems and illustrated its use. Second, on a more theoretical and conceptual level, we have uncovered a distinction between two types of dynamic inconsistency. We provide definitions for two types of consistency and show how they apply to complete and incomplete orderings.

Prior definitions of dynamic inconsistency tend to focus on one aspect: that a first-period optimal strategy can become nonoptimal in the second period. For complete preference orderings, this is equivalent to the converse: a first-period nonoptimal strategy becoming optimal in the second period. However, these two types of dynamic inconsistency are not equivalent for incomplete orders, such as dominance relations. We have explicitly defined these concepts, labelling them, respectively, fallacious and fallible inconsistency, and tied these to the definitions of weak and strong consistency.

When considering the merits of different decision rules, we argue that decision-makers face a trade-off between flexibility and full dynamic consistency. Subjective expected utility is dynamically consistent, but reduces decision-makers’ flexibility by narrowing in on a particular solution that hides the disagreement between beliefs. This can be problematic in situations involving multiple stakeholders, for example. Non-expected-utility decision rules are similarly inflexible and require specific recursive formulations to avoid fallacious inconsistency. Belief dominance is more consistent than nonrecursive non-expected-utility rules, avoiding fallacious inconsistency, but less consistent than expected utility, as it is subject to fallible inconsistency. However, by providing a range of defensible solutions, it leaves stakeholders the flexibility to negotiate and incorporate nonquantitative aspects into the decision-making process. It also can reduce the need to explicitly choose a particular weighting by providing a characterization of which weightings over beliefs lead to a given nondominated strategy, as we showed in the climate example.

To summarize, if significant disagreement exists over beliefs and the concern over fallible inconsistency is minor, then belief dominance provides a mechanism that eliminates bad alternatives while providing flexibility and preventing egregious errors of dynamic consistency. In some cases, it allows the best course of action to be identified without forcing stakeholders to agree on beliefs. On the other hand, if disagreement between beliefs is less important or if fallible inconsistency provides a great threat, then SEU should perhaps be used. The second case might arise if there are considerable commitment issues, for example, if one decision-maker is in charge of near-term decisions, while another will take over in the future.

One approach to narrowing down alternatives would be to apply a bottom-up exploratory approach, such as robust decision making (RDM; Rosenhead et al. 1972, Lempert and Collins 2007), decision scaling (Brown et al. 2012), or info gap (Ben-Haim 2004). These methods typically analyze a small set of predefined alternatives for robustness, then iterate to derive possible new alternatives. Many-objective robust decision making (MORDM) (Kasprzyk et al. 2013, Quinn et al. 2017) is somewhat parallel to our approach, but focuses on multiobjective problems and uses Pareto satisficing to derive a set of alternatives to analyze by RDM.

The dynamic setting, with new information gained between decisions, gives rise to possible updating of beliefs. Our formulation allows experts' individual beliefs to be subject to Bayesian updating between the periods. However, we do not apply Bayesian updating to the *set* of beliefs, since dominance does not use weighting or second-order probabilities over beliefs. In addition, it is worth noting that the beliefs considered in real-world decision-making settings can also be biased and time-inconsistent (Brunnermeier et al. 2017). Obvious extensions to our model include adding more periods and allowing the probabilities to depend on the earlier actions.

Appendix A: Formalization of the Problem from Section 2

Let Y be a random variable representing the virus strain, with $Y \in \{A, B\}$. The probabilities for $Y = A$ stated by the two virologists are presented in the top row of Table 1. Let $Z \in \{0,1\}^{12}$ be a 12-dimensional binary random variable indicating whether a specific vaccine is effective against strain A or strain B, accounting for whether one or two months of research have been put into the research of that vaccine. The probabilities for Z are presented in the bottom 12 rows of Table 1. Label the elements of Z as $Z_{Y,x}$, corresponding to research effort x being effective against strain Y.

Denote the number of effective vaccines, given strategy x, strain Y, and outcome Z with $N(x, Y, Z)$. Note that N may be zero, one, or even two if research efforts in two different vaccines are both successful against the strain Y. The expected value of $N(x, Y, Z)$ is given by

$$E_{Y,Z}[N(x,Y,Z)] = Pr(Y = A)\, E_Z[N(x,Y,Z)|Y = A] + Pr(Y = B)\, E_Z[N(x,Y,Z)|Y = B],$$

where $E[\cdot]$ stands for expectation. Further, $E_Z[N(x,Y,Z)|Y = y]$ equals $Pr(Z_{y,(x_1,x_{2,y})} = 1|Y = y)$ if $x_1 = x_{2,y}$ and $Pr(Z_{y,x_1} = 1|Y = y) + Pr(Z_{y,x_{2,y}} = 1|Y = y)$ if $x_1 \neq x_{2,y}$, with the probabilities as given in Table 1. For example, the expected value of strategy (β, α, β) under belief 1 is 90% (60% + 20%) + 10% · 80% = 0.8.

Appendix B: Proofs from Section 3

Proof for Theorem 1. *Case (ii):* $\tilde{x}_2(y) \sim_{\tilde{x}_1,y} x_2(y)$. Assume that weak consistency is violated: although $\tilde{x}_2(y) \sim_{\tilde{x}_1,y} x_2(y)$, it holds that $\tilde{x} \nsim x$. By strong consistency condition (c), since $\tilde{x} \sim \tilde{x}$ and $x_2(y) \sim_{\tilde{x}_1,y} \tilde{x}_2(y)$, we have that $(\tilde{x}_1, x_2(y), \tilde{x}_2(-y)) = x \sim \tilde{x}$. This contradicts the violation of weak consistency with $\tilde{x} \nsim x$.

Proof for Theorem 3. Assume that preference ordering $\preccurlyeq$ has the property of weak consistency, but fallacious inconsistency is observed. This means that $\exists x^* \in X^*$ with $x_2^*(y) \notin X^*_{2,x_1^*,y^*}$.

This implies that there is a strategy $\hat{x} = (x_1^*, \hat{x}_2, x_2^*(-y))$ with $\hat{x}_2 \succ_{x_1^*,y} x_2^*(y)$. Weak consistency implies that $\hat{x} \succ x$, but this contradicts the initial setting $x^* \in X^*$, as this means that $x^* \nprec \hat{x}$. Therefore, fallacious inconsistency cannot take place with weakly consistent preferences. ■

Proof for Theorem 4. Assume that preference ordering $\preccurlyeq$ has the property of strong consistency, but fallible inconsistency is observed. This means that

a. there is a solution $x^* \nprec x\ \forall x \in X$,
b. $\exists \hat{x}_2 \in X_{2,x_1^*,y}$ so that $\hat{x}_2 \nprec_{x_1^*,y} \check{x}_2$, $\forall \check{x}_2 \in X_{2,x_1^*,y}$,
c. $\exists \tilde{x} \in X$ so that $\tilde{x} \succ \hat{x} = (x_1^*, \hat{x}_2, x_2^*(-y))$.

Specifically, (b) above implies $\hat{x}_2 \nprec_{x_1^*,y} x_2^*(y)$. If strong consistency holds, then (b) and (a) together imply $\hat{x} = (x_1^*, \hat{x}_2, x_2^*(-y)) \nprec x$, $\forall x \in X$. This contradicts (c) above. Therefore, fallible inconsistency cannot occur with a preference ordering having strong consistency. ■

Endnotes

[1] We note that the terms uncertainty and risk carry different meanings in various academic communities. Knight (1921) differentiated risk to be measurable and uncertainty unmeasurable. The latter is sometimes called Knightian uncertainty or ambiguity, while Lempert et al. (2003) categorize this as one class of "deep uncertainty." In some communities—in risk analysis, for example—risk is defined as a possibility of loss, while uncertainty can be quantified through probabilities (Kaplan and Garrick 1981). In this paper, we use uncertainty as a general term for the lack of certainty, risk for uncertainty described by probabilities, and ambiguity for unknown probabilities, similar to, for example, Bertsimas et al. (2019).

[2] Although minmax regret violates the independence of irrelevant alternatives (see, e.g., Stoye 2012a), a basic axiom of rational choice, it is widely used in the literature (e.g., Caldentey et al. 2017, van der Pol et al. 2017, Xidonas et al. 2017).

[3] This illustrative example was conceived in early 2018, that is, it predates the COVID-19 pandemic.

[4] We assume here for simplicity that carrying out research for the first month will not yield better estimates for the probabilities, which the manager could then use to inform the second-period decision.

[5] That is, in the terms of order theory, the relation $\succcurlyeq$ is not a total preorder.

[6] We wish to thank an anonymous referee for making this observation.

[7] The reader can refer to, for example, Høyland and Wallace (2001) for a method for discretizing continuous probabilities for scenario trees.

References

Al-Najjar NI, Weinstein J (2009) Rejoinder: The ambiguity aversion literature: A critical assessment. *Econom. Philos.* 25:357–369.
Aumann RJ (1962) Utility theory without the completeness axiom. *Econometrica* 30(3):445–462.
Baker E, Bosetti V, Salo A (2020) Robust portfolio decision analysis: An application to the energy research and development portfolio problem. *European J. Oper. Res.* 284(3):1107–1120.
Battigalli P, Catonini E, Lanzani G, Marinacci M (2019) Ambiguity attitudes and self-confirming equilibrium in sequential games. *Games Econom. Behav.* 115:1–29.
Bell DE (1982) Regret in decision making under uncertainty. *Oper. Res.* 30:961–981.
Ben-Haim Y (2004) Uncertainty, probability and information-gaps. *Reliability Engrg. System Safety* 85:249–266.
Ben-Tal A, den Hertog D, De Waegenaere A, Melenberg B, Rennen G (2013) Robust solutions of optimization problems affected by uncertain probabilities. *Management Sci.* 59:341–357.
Bertsimas D, Sim M, Zhang M (2019) Adaptive distributionally robust optimization. *Management Sci.* 65:604–618.
Bewley TF (2002) Knightian decision theory. Part I. *Decisions Econom. Finance* 25:79–110.

Brown C, Ghile Y, Laverty M, Li K (2012) Decision scaling: Linking bottom-up vulnerability analysis with climate projections in the water sector. *Water Resources Res.* 48(9):W09537.

Brunnermeier MK, Papakonstantinou F, Parker JA (2017) Optimal time-inconsistent beliefs: Misplanning, procrastination, and commitment. *Management Sci.* 63:1318–1340.

Caldentey R, Liu Y, Lobel I (2017) Intertemporal pricing under minimax regret. *Oper. Res.* 65:104–129.

Cox LA (2012) Confronting deep uncertainties in risk analysis. *Risk Anal.* 32:1607–1629.

Danan E, Gajdos T, Hill B, Tallon J-M (2016) Robust social decisions. *Amer. Econom. Rev.* 106:2407–2425.

Diaz D, Moore F (2017) Quantifying the economic risks of climate change. *Nature Climate Change* 7:774–782.

Ekholm T (2018) Climatic cost-benefit analysis under uncertainty and learning on climate sensitivity and damages. *Ecological Econom.* 154:99–106.

Epstein LG, Schneider M (2003) Recursive multiple-priors. *J. Econom. Theory* 113:1–31.

Etner J, Jeleva M, Tallon JM (2012) Decision theory under ambiguity. *J. Econom. Surveys* 26:234–270.

Geoffrion AM (1968) Proper efficiency and the theory of vector maximization. *J. Math. Anal. Appl.* 22:618–630.

Ghirardato P, Maccheroni F, Marinacci M (2004) Differentiating ambiguity and ambiguity attitude. *J. Econom. Theory* 118:133–173.

Ghirardato P, Maccheroni F, Marinacci M (2008) Revealed ambiguity and its consequences: Updating. Abdellaoui M, Hey JD, eds. *Advances in Decision Making Under Risk and Uncertainty* (Springer, Berlin, Heidelberg), 3–18.

Gilboa I, Schmeidler D (1989) Maxmin expected utility with non-unique prior. *J. Math. Econom.* 18:141–153.

Gilboa I, Marinacci M (2016) Ambiguity and the Bayesian paradigm. Arló-Costa H, Hendricks VF, van Benthem J, eds. *Readings in Formal Epistemology* (Springer International Publishing, Switzerland), 385–439.

Gilboa I, Maccheroni F, Marinacci M, Schmeidler D (2010) Objective and subjective rationality in a multiple prior model. *Econometrica* 78:755–770.

Grushka-Cockayne Y, De Reyck B, Degraeve Z (2008) An integrated decision-making approach for improving European air traffic management. *Management Sci.* 54:1395–1409.

Gul F, Pesendorfer W (2001) Temptation and self-control. *Econometrica* 69:1403–1435.

Hayashi T (2005) Intertemporal substitution, risk aversion and ambiguity aversion. *Econom. Theory* 25:933–956.

Hill B (2020) Dynamic consistency and ambiguity: A reappraisal. *Games Econom. Behav.* 120:289–310.

Høyland K, Wallace SW (2001) Generating scenario trees for multistage decision problems. *Management Sci.* 47:295–307.

Johnson MP, Hurter AP (2000) Decision support for a housing mobility program using a multiobjective optimization model. *Management Sci.* 46:1569–1584.

Kaplan S, Garrick BJ (1981) On the quantitative definition of risk. *Risk Anal.* 1:11–27.

Kasprzyk JR, Nataraj S, Reed PM, Lempert RJ (2013) Many objective robust decision making for complex environmental systems undergoing change. *Environ. Model. Software* 42:55–71.

Klibanoff P, Marinacci M, Mukerji S (2005) A smooth model of decision making under ambiguity. *Econometrica* 73:1849–1892.

Klibanoff P, Marinacci M, Mukerji S (2009) Recursive smooth ambiguity preferences. *J. Econom. Theory* 144:930–976.

Knight F (1921) *Risk, Uncertainty and Profit, Risk* (Augustus M. Kelley, New York).

Knutti R, Hegerl GC (2008) The equilibrium sensitivity of the Earth's temperature to radiation changes. *Nature Geoscience* 1:735–743.

Lemoine D, Rudik I (2017) Managing climate change under uncertainty: Recursive integrated assessment at an inflection point. *Annu. Rev. Resources Econom.* 9:117–142.

Lempert R, Collins M (2007) Managing the risk of uncertain threshold responses: Comparison of robust, optimum, and precautionary approaches. *Risk Anal.* 27(4):1009–1026.

Lempert RJ, Popper SW, Bankes SC (2003) *Shaping the next one hundred years: New methods for quantitative long-term strategy analysis* (RAND, Santa Monica, CA).

Loomes G, Sugden R (1982) Regret theory: An alternative theory of rational choice under uncertainty. *Econom. J. (London)* 92: 805–824.

Maccheroni F, Marinacci M, Rustichini A (2006) Dynamic variational preferences. *J. Econom. Theory* 128:4–44.

Machina MJ (1989) Dynamic consistency and non-expected utility models of choice under uncertainty. *J. Econom. Lit.* 27:1622–1668.

Marinacci M (2015) Model uncertainty. *J. Eur. Econom. Assoc.* 13: 1022–1100.

Millner A, Dietz S, Heal G (2013) Scientific ambiguity and climate policy. *Environ. Resources Econom.* 55:21–46.

Morton A (2014) Aversion to health inequalities in healthcare prioritisation: A multicriteria optimisation perspective. *J. Health Econom.* 36:164–173.

Nordhaus WD (1992) An optimal transition path for controlling greenhouse gases. *Science* 258:1315–1319.

Nordhaus WD (2017) Revisiting the social cost of carbon. *Proc. Natl. Acad. Sci. USA* 114:1518–1523.

Pindyck RS (2013) Climate change policy: What do the models tell us? *J. Econom. Lit.* 51:860–872.

Quinn JD, Reed PM, Keller K (2017) Direct policy search for robust multi-objective management of deeply uncertain socio-ecological tipping points. *Environ. Model. Software* 92:125–141.

Rosenhead J, Elton M, Gupta SK (1972) Robustness and optimality as criteria for strategic decisions. *Oper. Res. Quart.* 23(4):413–431.

Savage LJ (1954) *The Foundations of Statistics* (Dover Publications, New York).

Siniscalchi M (2009) Two out of three ain't bad: A comment on the ambiguity aversion literature: a critical assessment. *Econom. & Philos*. 25(3):335–356.

Siniscalchi M (2011) Dynamic choice under ambiguity. *Theoret. Econom*. 6(3):379–421.

Stoye J (2012a) New perspectives on statistical decisions under ambiguity. *Annu. Rev. Econom.* 4:257–282.

Stoye J (2012b) Dominance and admissibility without priors. *Econom. Lett.* 116:118–120.

Strotz RH (1955) Myopia and inconsistency in dynamic utility maximization. *Rev. Econom. Stud.* 23:165–180.

Thaler RH, Shefrin HM (1981) An economic theory of self-control. *J. Political Econom.* 89:392–406.

van der Pol TD, Gabbert S, Weikard HP, van Ierland EC, Hendrix EMT (2017) A minimax regret analysis of flood risk management strategies under climate change uncertainty and emerging information. *Environ. Resources Econom.* 68:1087–1109.

Wald A (1949) Statistical decision functions. *Ann. Math. Statist.* 20:165–205.

Wang T (2003) Conditional preferences and updating. *J. Econom. Theory* 108:286–321.

Wiesemann W, Kuhn D, Sim M (2014) Distributionally robust convex optimization. *Oper. Res.* 62:1358–1376.

Xidonas P, Mavrotas G, Hassapis C, Zopounidis C (2017) Robust multiobjective portfolio optimization: A minimax regret approach. *Eur. J. Oper. Res.* 262:299–305.