# TESTING DYNAMIC CONSISTENCY AND CONSEQUENTIALISM UNDER AMBIGUITY

HAN BLEICHRODT, JÜRGEN EICHBERGER, SIMON GRANT, DAVID KELSEY, AND CHEN LI

ABSTRACT. Accounting for ambiguity aversion in dynamic decisions generally implies that either dynamic consistency or consequentialism must be given up. To gain insight into which of these principles better describes people's preferences, we tested them using a variation of Ellsberg's three-color urn experiment. Subjects were asked to make a choice both before and after they received a signal. We found that most ambiguity neutral subjects satisfied both dynamic consistency and consequentialism and behaved consistent with subjective expected utility with Bayesian updating. The majority of ambiguity averse subjects satisfied consequentialism, but violated dynamic consistency.

KEYWORDS: ambiguity, three-color Ellsberg paradox, consequentialism, dynamic consistency.

JEL CLASSIFICATION. C72, D81

## 1. INTRODUCTION

Many real-world decisions involve ambiguity, where outcomes are uncertain and their probabilities unknown. Moreover, most of these decisions are dynamic in the sense that information becomes available over time, probabilities are updated, and optimal strategies or policies may be revised. Examples are the threats from a new disease and climate change. The classic model to analyze such decisions is subjective expected utility, which assumes that decision makers are ambiguity neutral, with updating of probabilities according to Bayes' rule. However, since Ellsberg (1961), empirical research has shown that people

*Date*: 27th January 2021.

We are grateful to editors Steffen Huck and David Levine, an associate editor, and 4 reviewers for their helpful suggestions on previous versions of our paper. The authors would like to thank the Erasmus Research Institute of Management and the Erasmus School of Economics for funding the experiment reported in this paper.

do not behave according to subjective expected utility, but are averse to ambiguity, and update in non-Bayesian manners.

There are by now a number of well-established theories to capture such ambiguity aversion, but they are essentially static in nature.[1] To extend these theories to dynamic decisions raises an interesting problem. Several authors have shown that modeling ambiguity aversion in dynamic decisions implies giving up one of two principles that are usually considered rational: *consequentialism*, the requirement that preferences only depend on outcomes on states that are still possible, and *dynamic consistency*, the requirement that decision makers stick to their optimal contingent plans.[2] Subjective expected utility maximizers who update their beliefs according to Bayes' rule satisfy both consequentialism and dynamic consistency, but ambiguity averse decision makers cannot. Hence, dynamic ambiguity models must choose which of these principles to retain.[3]

The purpose of our paper is to inform this choice by providing experimental evidence on the descriptive validity of these two principles. The question which of these principles has more descriptive appeal is also of

[1] In these models beliefs either cannot be represented by a (single) probability measure or the probability measure is a second-order belief and the corresponding induced preferences do not satisfy reduction of compound lotteries. Examples of the former are the multiple priors model of Gilboa and Schmeidler (1989) and the Choquet expected utility model of Schmeidler (1989). Examples of the latter include the two-stage lotteries without reduction model of Segal (1990), the second-order beliefs models of Klibanoff, Marinacci, and Mukerji (2005), Nau (2006) and Seo (2009), and the subjective compound lottery model of Ergin and Gul (2009).

[2] Ghirardato (2002a) showed this in a Savage framework without the sure-thing principle and with a more intuitive version of Savage's axiom P3. Epstein and Le Breton (1993) showed in a framework without Savage's P2 that dynamic consistency implies probabilistic sophistication (Machina and Schmeidler, 1992). See also Siniscalchi (2011).

[3] Machina (1989) argued to give up consequentialism. Studies that followed his recommendation are Hanany and Klibanoff (2007), Hanany and Klibanoff (2009), and Klibanoff, Marinacci, and Mukerji (2009).

Examples of studies that give up dynamic consistency and retain consequentialism are Karni and Safra (1990), Gilboa and Schmeidler (1993), and Siniscalchi (2011). Dropping dynamic consistency is also common in the literature on hyperbolic discounting (O'Donoghue and Rabin, 1999).

interest for applied economic research (standard dynamic optimization techniques require dynamic consistency) and for policy. For example, consequentialism may need to be given up if policy makers care not only about the outcome of decisions, but also about the process by which they are arrived at. This has been emphasized for example in the provision of health care.

Our experiment is a variation of the Ellsberg three-color urn problem Ellsberg (1961). We used 200 cards that were both numbered and colored. The color composition was the same for even- and odd-numbered cards: for each parity there were 33 red cards while the remaining 67 cards were blue and yellow in unknown proportions. Subjects chose a color to bet on for both odd and even numbered cards. This allowed them to hedge against ambiguity. For instance, if a subject chose to bet on blue for odd cards and yellow for even cards then, even though they did not know the exact number of blue odd cards or yellow even cards, they knew that this combination gave a 67/200 chance to win. Subjects then received a signal whether the winning card was odd or even. This signal removed the possibility to hedge against ambiguity: conditional on knowing whether the cards was even or odd, the choice became similar to the standard Ellsberg three-color problem.

By comparing subjects' choices before and after the signal, we could test dynamic consistency and consequentialism. An ambiguity averse subject who satisfied dynamic consistency would make the same choice before and after the signal, whereas an ambiguity averse subject who satisfied consequentialism would switch choices before and after the signal. To distinguish ambiguity averse subjects clearly and to detect violations of dynamic consistency, we included a small cost for being ambiguity averse. This also allowed us to test to what extent subjects' preferences were monotonic in probability and satisfied (first-order) stochastic dominance.

We observed that 42% of our subjects were ambiguity neutral. Around two thirds of these ambiguity neutral subjects satisfied both consequentialism and dynamic consistency and behaved consistent with subjective expected utility with Bayesian updating. Around one third of our subjects were ambiguity averse. Most of these behaved according to consequentialism and violated dynamic consistency: violations of dynamic consistency occurred about twice as often as those of consequentialism. This provides support for dropping dynamic consistency in dynamic decision problems when accounting for ambiguity aversion.

Dominiak, Duersch, and Lefort (2012) also tested dynamic consistency and consequentialism in a dynamic Ellsberg experiment.[4] Like us, they also found that more violations of dynamic consistency than consequentialism. In fact, the proportion of subjects satisfying consequentialism in their study was the same as the proportion we observed (73%). Their test was not based on actual information, but they asked their subjects to imagine that some information was revealed to them. By contrast, in our experiment subjects made their choices based on the actual information that was revealed to them. Another difference between our study and Dominiak, Duersch, and Lefort (2012) is the treatment of indifference. They used unincentivised verbal statements about strength of preference to distinguish between ambiguity averse and ambiguity neutral subjects. We derive this distinction from incentivised choices by incurring a small cost to be ambiguity averse. This small cost might be responsible for the lower proportion of ambiguity averse subjects in our study. It is reassuring that in spite of these

[4]A related paper on decision-making under risk is Cubitt, Starmer, and Sugden (1998).

differences in design our conclusions are similar: if given the choice, ambiguity averse subjects are more likely to satisfy consequentialism than dynamic consistency.

The following Section 2 presents the design of the experiment. Section 3 defines consequentialism and dynamic consistency and it explains how our experiment tested them. Section 4 presents our results. Section 5 discusses the results in the context of the literature. Section 6 concludes.

## 2. Experiment

We conducted a lab experiment to elicit subjects' ambiguity attitudes and to test whether they satisfied consequentialism and dynamic consistency. Screenshots of the experimental instructions and all decision situations are in the online appendix.[5]

**Subjects and Incentives.** The experiment was computer-run and conducted in the ESE-Econlab at Erasmus University Rotterdam in March 2018. The experiment consisted of 7 sessions, with 23 to 27 subjects per session. In each session, two subjects were randomly assigned to be implementers. Their role was to generate ambiguity by determining the color composition of bags with cards, as explained below. A total of 171 subjects were recruited from the ESE-Econlab subject pool of whom 14 were assigned to be implementers. Data were collected from the remaining 157 subjects.

The experiment was incentivized using the prior incentive system (Prince; Johnson, Baillon, Bleichrodt, Li, van Dolder, and Wakker 2020). Upon entering the lab, every subject drew an envelope from

[5] https://www.dropbox.com/s/2d5wr5fg7vb1y0i/online-appendix.pdf?dl=0

a pile of $n$ sealed envelopes ($n$ = number of participants in each session). A subject ID was written on each envelope. Subjects who drew an ID starting with "m" became the implementers.

We told the subjects in the main experiment that their envelope contained one of the decision situations they would encounter during the experiment. Each decision situation during the experiment has the same chance of being drawn. Subjects could only open their envelopes at the end of the experiment when the experimenters told them so. They would be paid for real according to their choice in the decision situation contained in their envelope, which varied across subjects. The implementers received a flat participation fee of €10. The subjects in the main experiment received a participation fee of €5 and a variable amount, which depended on their choices in the decision situation in their envelope. The average payment was €9.62.

**Stimuli.** Subjects faced four decision situations. The first two situations measured subjects' ambiguity attitudes and were based on Ellsberg's 3-color problem using a bag containing 100 *unnumbered* cards colored red, blue, or yellow. The final two decision situations were based on the draw of a card from a second bag containing 200 cards that were both colored (red, blue, or yellow) and *numbered.* Per experimental session, there were two implementers, one to determine the color composition of the cards in the first bag and one to determine (independently) the color composition of the cards in the second bag. The implementers did not know how the composition they determined might affect the other subjects' payoffs and the subjects were told that the implementers did not know this

**Decision situations 1 and 2: Test of ambiguity attitude:**

The first implementer was asked to write down a number $N$ between 1 and 67 and then to put 33 red cards, $N$ blue cards, and $67-N$ yellow cards into a bag. Subjects knew this procedure but they did not know the number $N$ that the implementer had written down. After having filled the bag, the implementer drew the winning card.

Subjects could bet on their winning colors. If a card with a winning color was drawn, subjects won €10. Otherwise, they won nothing. In decision situation 1, subjects could bet on red, blue, or yellow. In the second decision situation, subjects could bet on two colors depending on which winning color they had selected in decision situation 1. If they had chosen red or blue in decision situation 1, then they were offered a choice between "red and yellow" and "blue and yellow". If they had chosen yellow in decision situation 1, then they were offered a choice between "red and blue" and "blue and yellow". We let subjects choose their wining color to avoid suspicion.

In decision situation 1, "red" was unambiguous and "blue" and "yellow" were ambiguous. In decision situation 2, "red and yellow" and "red and blue" were ambiguous whereas "blue and yellow" was unambiguous. Consequently, we expected an ambiguity averse decision maker to choose "red" in the first choice and "blue and yellow" in the second choice. The expected probability of drawing a blue or yellow card in decision situation 1 was 0.335, which was slightly higher than the 0.33 probability of drawing a red card. Hence, the expected loss of ambiguity aversion was 5 cents. We put a small price on ambiguity aversion to avoid indifference. Kelsey and LeRoux (2017) showed that even small prices can substantially affect the number of ambiguity averse choices. A limitation of our approach was that some (very weakly) ambiguity averse decision makers might choose $B$ or $Y$. Consequently, the proportion of ambiguity averse choices that we observed may be a lower estimate. Given the low price of 5 cents, this downward bias is probably small to negligible.

**Decision situations 3 and 4: dynamic decisions**.

The second implementer wrote down a number $M$ between 1 and 67 and then put 66 red cards (numbered from 135 to 200), $2M$ blue cards (numbered from 1 to $2M$), and 134-$2M$ yellow cards (numbered from $2M$+1 to 134) into a bag. Again, the subjects were aware of this procedure, but they did not know the number $M$ that the second implementer had written down. After having filled the bag, the implementer drew the winning card.

Subjects first answered three comprehension questions. They could only proceed to the main experiment after having answered all of these correctly.

In decision situation 3, subjects had to specify their winning color for both odd- and even-numbered cards.

After they had made their choices in decision situation 3, implementer 2 revealed the parity of the number on the card that they had drawn. Then in decision situation 4, subjects indicated which color they would like to bet on given this signal about the parity of the card drawn.

At the end of the experiment, the two implementers revealed the color and the number of the drawn cards to the other participants.

The structure of the experiment is summarized in figure 1.

## 3. Theory: Consequentialism versus dynamic inconsistency

This section presents our theoretical analysis of subjects' choices in the experiment. Uncertainty is modeled by a state space $S$. Single states are denoted $s_1, s_2, ....$ An *event* $E$ is a subset of the set $S$. In

Figure 1. **Structure of the Experiment**

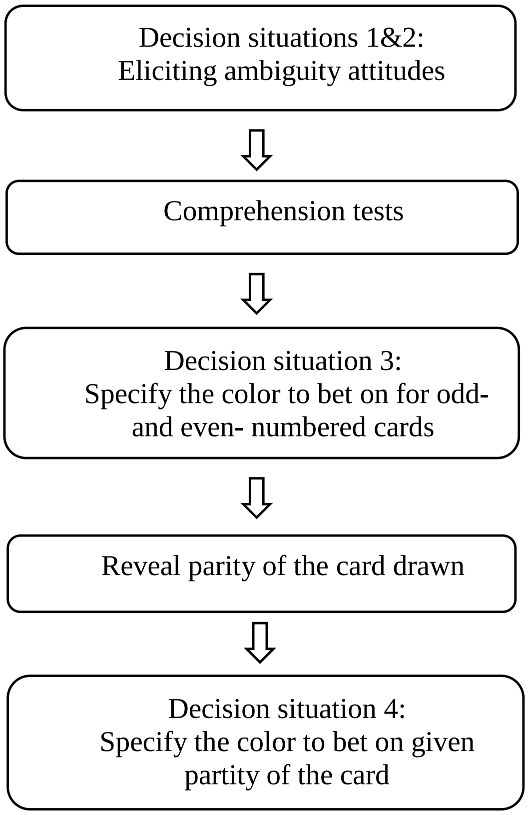


our experiment, the events in decision situations 1 and 2 are drawing a red, blue, or yellow card, and in decision situations 3 and 4 they are all combinations of color and parity of the card, such as, *drawing a red card with an even number on it* or *drawing a yellow card with an odd number on it.* We denote the complement of an event $E$ by $E^c$. A *strategy* defines at each decision node what to do. For example, in decision situation 3 the strategy $Y_o R_e$ stands for *bet on yellow when the number on the card is odd, bet on red when it is even.* We follow Sarin and Wakker (1998), by defining preferences $\succsim$ over the set of strategies.

Figure 2 shows the decision situations of our experiment that are relevant for testing consequentialism and dynamic consistency.[6] Squares

[6]Decision situation 2 was only relevant for testing ambiguity aversion and we, therefore, do not display it in the figure.

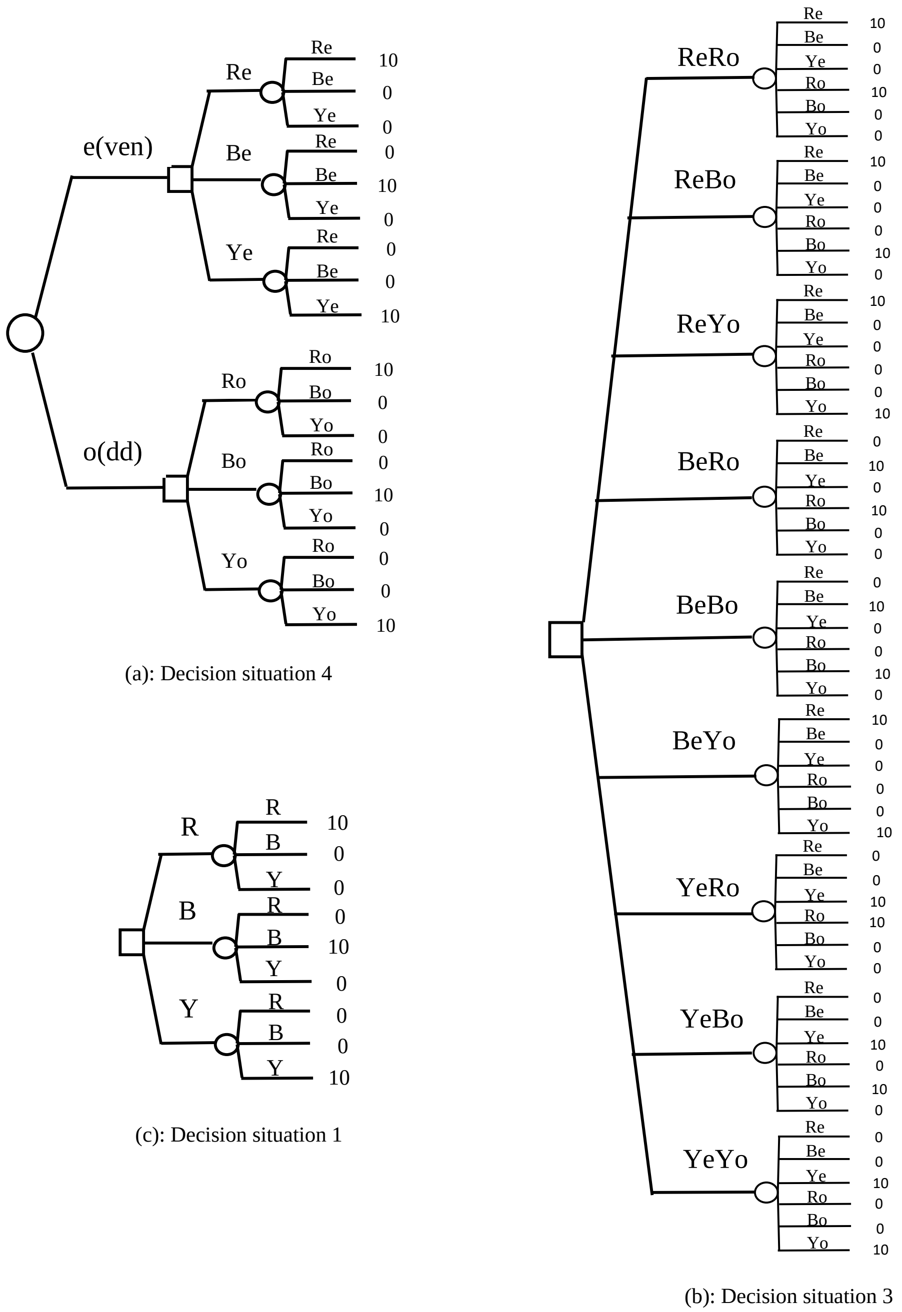


Figure 2. Representation of decision situations in our experiment. Squares represent choice nodes and circles represent chance nodes.

denote decision nodes and circles denote chance nodes. Figure 2(b) shows decision situation 3, which is a *single-stage decision* in which decision makers must commit beforehand to their choices, that is before the uncertainty about the parity of the winning card is resolved. Figure 2(a) shows decision situation 4, which is an example of a *dynamic decision*. In decision situation 3, subjects had specified on which color to bet when the number were odd and when it were even. They now receive information whether the winning card has an odd or an even number on it and are asked to make another choice. These decisions are made after the uncertainty about the parity of the card has been resolved. Finally, Figure 2(c) shows decision situation 1, which can be interpreted as a *no-history decision,* and which is similar to the upper or lower branch of Figure 2(a) once the parity of the card is known.

*Dynamic consistency* is the requirement that choices in the dynamic decision agree with those in the single-stage decision. So if strategy $Y_oR_e$ is chosen in (the single-stage) decision situation 3 (Figure 2(b)) then in the (dynamic) decision situation 4, $Y$ must be chosen if the number on the drawn card is odd and $R$ if the number on the card is even. In other words, dynamic consistency requires that decision makers do not change their minds after the uncertainty resolves at the chance nodes.

By the design of our experiment, we gave subjects real information (whether the number on the drawn card was even or odd). Hence, we could only observe subjects' conditional choices for the event that actually occurred and not those for the counterfactual event. For instance, when an odd-numbered card was drawn, we could only observe subjects' choices in the lower branch of Figure 2(a), but not in the upper branch. Therefore, at the individual subject level, we did not have a full test of dynamic consistency. Nevertheless, since the two conditional decision situations (for odd- and even-numbered cards) are informationally indistinguishable, we believe that it is reasonable to assume that subjects' choices would not differ in these two situations. Our data support this assumption. In the aggregate, roughly half of the subjects faced each conditional bet and the color choices did not differ between the case where the card was even and the case where it was odd. Statistical tests could not reject the null that the choice distributions were equal in the two conditional decision situations.[7]

*Consequentialism* means that the choices made in dynamic decision situations are independent of risks born in the past, see Hammond (1988).[8] In our experiment, it means that the choices in the upper and lower branches of Figure 2(a) agree with those in the no-history decision situation 1, which isolates these decisions. Consequentialism requires that choices made after the resolution of uncertainty are not affected by what would have happened in the counterfactual event. Hence, choices are history-independent, and coincide with those made in "no-history" decision situations as depicted in Figure 2(c). Comparing the choice in Figure 2(c) with that in Figure 2(a) constitutes our test of consequentialism. This test requires the additional assumption that subjects perceived the two decision situations as similar even though the uncertainty was resolved by different implementers. If subjects did not consider the two decision situations as similar, for instance because they held different beliefs about the implementers, then they might fail our test of consequentialism even though their underlying preferences satisfied consequentialism and, vice versa, they might pass our test of

[7]In decision situation 4, of the decision makers who made an "odd-numbered contingent" choice, 51% preferred red, 30% preferred blue and 19% preferred yellow, whereas for those who made an "even-numbered contingent" choice, 46% preferred red, 30% preferred blue and 24% preferred yellow ($p = 0.68$ in Fisher's Exact test).

[8]A more technical definition of consequentialism is included in the appendix.

consequentialism even though their underlying preference relation did not satisfy consequentialism. Nevertheless, we believe that the assumption is plausible given that the two conditional decision situations and the no-history decision situation are informationally indistinguishable and that subjects typically did not know the implementers. The similarity between the violation rates of consequentialism in our experiment and those in Dominiak, Duersch, and Lefort (2012) lends empirical credence to our test.

Consequentialist decision makers with non-neutral attitudes toward ambiguity may be dynamically inconsistent in our problems as the partial resolution of uncertainty (revealing whether the drawn card had an even or an odd number on it) changes ambiguous strategies into *conditionally* unambiguous strategies and vice versa. To illustrate this observation, consider the strategies in 2(b). Given the composition of the bag, decision-makers know the strategy $R_oR_e$ has 66 winning cards (33 odd-numbered red cards and 33 even-numbered red cards). Thus choosing $R_oR_e$ amounts to selecting the unambiguous strategy $10_{R_oR_e}0$ (which stands for winning €10 if the selected card is red and odd or red and even and nothing otherwise) with $p(R_oR_e) = 66/200$. The strategy $B_oY_e$ has 67 winning cards ($M$ odd-numbered blue cards and $67 - M$ even-numbered yellow cards). Thus the strategy $B_oY_e$ corresponds to the unambiguous strategy $10_{B_oY_e}0$ with $p(B_oY_e) = 67/200$.

By stochastic dominance, we expect decision makers to prefer $B_oY_e$ to $R_oR_e$. Similarly, we also expect them to prefer $Y_oB_e$ to $R_oR_e$. However, the information about the parity of the number on the card drawn changes the ambiguity of some of the options. The originally undominated unambiguous strategies $B_oY_e$ and $Y_oB_e$ become ambiguous, whereas the originally dominated unambiguous strategy $R_oR_e$ remains unambiguous.

Consequentialist decision makers who are ambiguity seeking will initially choose either $B_oB_e$ or $Y_oY_e$. Without loss of generality, suppose that they initially select $B_oB_e$. Regardless of the parity of the number on the card drawn, they will strictly prefer the ambiguous bet on blue over the unambiguous bet on red and weakly prefer it over the bet on yellow. So there is no reason for ambiguity seeking consequentialists to reverse their initial preference.

For ambiguity averse consequentialists, however, this is different. They initially choose one of the undominated unambiguous strategies $B_oY_e$ and $Y_oB_e$. Suppose that they initially choose $B_oY_e$ and that they are told that the number on the card drawn is odd. Since consequentialists no longer pay attention to what would have happened for even-numbered cards, they will be (conditionally) indifferent between $R_oR_e$ and $R_oY_e$. Moreover, if they are sufficiently ambiguity averse then they will (strictly) prefer a bet on red over a bet on blue. Thus, they will conditionally prefer $R_oY_e$ to $B_oY_e$. By transitivity, it follows that $R_oR_e$ is strictly preferred to $B_oY_e$ *reversing* their initial preference of $B_oY_e$ over $R_oR_e$. By an analogous argument, they will prefer $R_oR_e$ to $B_oY_e$ when they learn that the drawn ball is even.

The behavior of ambiguity seeking dynamically consistent decision makers will be indistinguishable from that of ambiguity loving consequentialists. They will initially select either $B_oB_e$ or $Y_oY_e$ and have no reason to change their choice after being informed about the parity of the number of the card drawn.

Ambiguity averse dynamically consistent decision makers, however, unconditionally prefer $B_oY_e$ over $R_oR_e$ and $R_oR_e$ over $R_oY_e$ and so, by transitivity, they prefer $B_oY_e$ over $R_oY_e$. If they know that an odd-numbered ball has been drawn then, by dynamic consistency, they will prefer betting on blue over red. By a similar line of reasoning, if they know that an even-numbered ball has been drawn they will prefer betting on yellow over red.

Consequently, we have derived the following observation.

**Observation 1.** *If stochastic dominance holds then ambiguity averse decision makers whose preferences satisfy*

- *consequentialism will violate dynamic consistency by choosing* $B_oY_e$ *over* $R_oR_e$ *in decision situation 3 (Figure 2b) and red over blue in decision situation 4 (Figure 2a).*
- *dynamic consistency will violate consequentialism by choosing* $B_oY_e$ *over* $R_oR_e$ *in decision situation 3 (Figure 2b) and blue over red in decision situation 3 (Figure 2a).*

The following picture summarizes the tests in our experiment.

**Tests of the Experiment:**

(1) Ambiguity attitude: decision situations 1 and 2
(2) Consequentialism: decision situations 1 and 4
(3) Dynamic Consistency: decision situations 3 and 4
(4) Stochastic Dominance: decision situation 3

## 4. Results

We report the results using the responses from all subjects.[9]

**Ambiguity attitudes.** We classified subjects who chose red in decision situation 1 and blue and yellow in decision situation 2 as ambiguity averse and subjects who chose blue or yellow in decision situation 1 and red and yellow or red and blue in decision situation 2 as ambiguity seeking. Because a small price to be ambiguity averse and the expected probability of blue or yellow card was slightly higher than that of red, an ambiguity neutral subject would choose blue or yellow in decision

[9]We performed several robustness tests. We also analyzed the data removing the subjects who (i) violated stochastic dominance; (ii) failed the comprehension tests at least twice. This did not affect our conclusions. Details are in the online appendix.

situation 1 and blue and yellow in decision situation 2. The remaining subjects (who chose red in decision situation 1 and red and yellow in decision situation 2) were classified as mixed.

Figure 3 shows that ambiguity neutrality was the modal pattern, followed by ambiguity aversion. Few subjects were ambiguity seeking. Our finding of limited ambiguity aversion is not uncommon, especially if subjects have to pay a price to pay to be ambiguity averse (see Trautmann and van de Kuilen, 2015 for a review of the empirical literature). In addition, it might be that some of the subjects classified as mixed were actually ambiguity averse, but failed to notice that blue and yellow in decision situation 2 was unambiguous.[10]

**Dynamic Decisions.** Figure 4 breaks down the choice pattern in decision situation 3 by ambiguity attitude. The proportions above the bars show for each category which share displayed which choice pattern.

Among the ambiguity averse subjects, 80% chose one of the three unambiguous options ($R_oR_e$, $B_oY_e$ and $Y_oB_e$). This is more than the proportion (41%) among ambiguity neutral subjects (two-tailed proportion test, $p < 0.01$, $N_A = 51$, $N_N = 66$). However, it is not different from the proportion among subjects classified as "mixed" (two-tailed proportion test, $p = 0.33$, $N_A = 51$, $N_B = 28$).

Forty-three subjects (27.4%) chose $R_oR_e$ over $B_oY_e$ and $Y_oB_e$, violating stochastic dominance. Most of these subjects were ambiguity averse or mixed. Relatively few ambiguity neutral subjects chose $R_oR_e$.[11] This proportion may appear high, but it should be kept in

[10]Their answers in decision situation 3 provided support for this conjecture. 70% of the mixed subjects chose one of the ambiguity averse options in decision situation 3 ($R_oR_e$,$B_oY_e$,or $Y_oB_e$). We found no evidence that the mixed subjects had more difficulty understanding the tasks and that their answers reflected confusion. There was no relation between the number of failures in the comprehension tests and ambiguity attitude ($\chi^2$-test, $p = 0.29$).

[11]Excluding the subjects who violated stochastic dominance did not affect our main conclusions. However it led to an increase in support for subjective expected

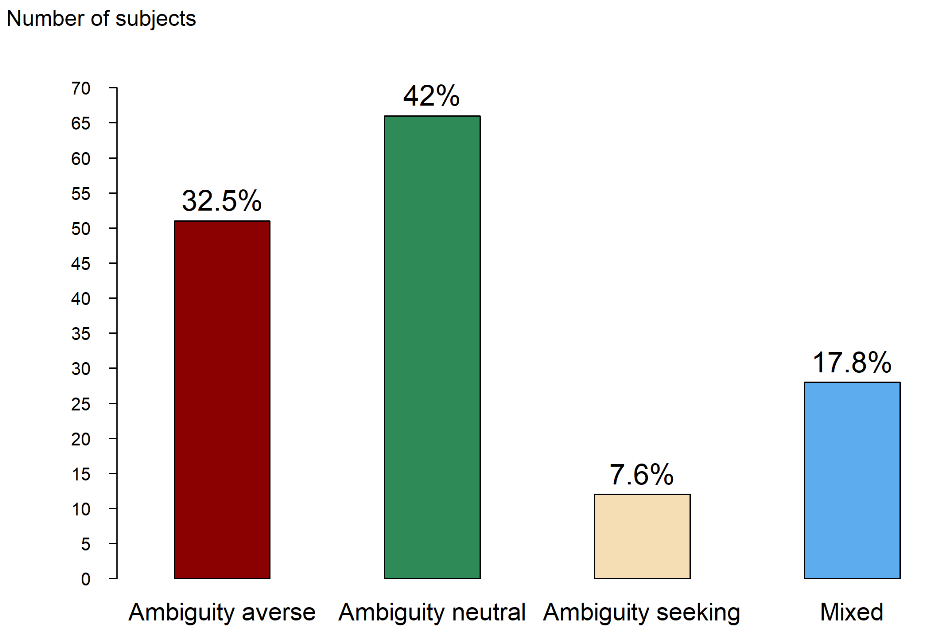


FIGURE 3. Ambiguity Attitudes

mind that the cost of violating stochastic dominance was low. It was also easier to detect that $R_oR_e$was unambiguous than that $B_oY_e$ and $Y_oB_e$ were unambiguous by providing a hedge against ambiguity. Some subjects may have considered the cost of 5 cents of violating stochastic dominance too low to take the cognitive effort of understanding why $B_oY_e$ and $Y_oB_e$ provided a hedge against ambiguity.[12] We discuss the violations of stochastic dominance in Section 5.

Figure 5 shows how many subjects behaved in line with dynamic consistency and consequentialism, split out by ambiguity attitudes. In the figure's legend, $DC$ and $C$ correspond to a subject satisfying the

utility with Bayesian updating: after the exclusion 35.2% of the subjects behaved according to subjective expected utility with Bayesian updating. See the online appendix for details.

[12]On the other hand, the third comprehension test explained how to hedge against ambiguity.

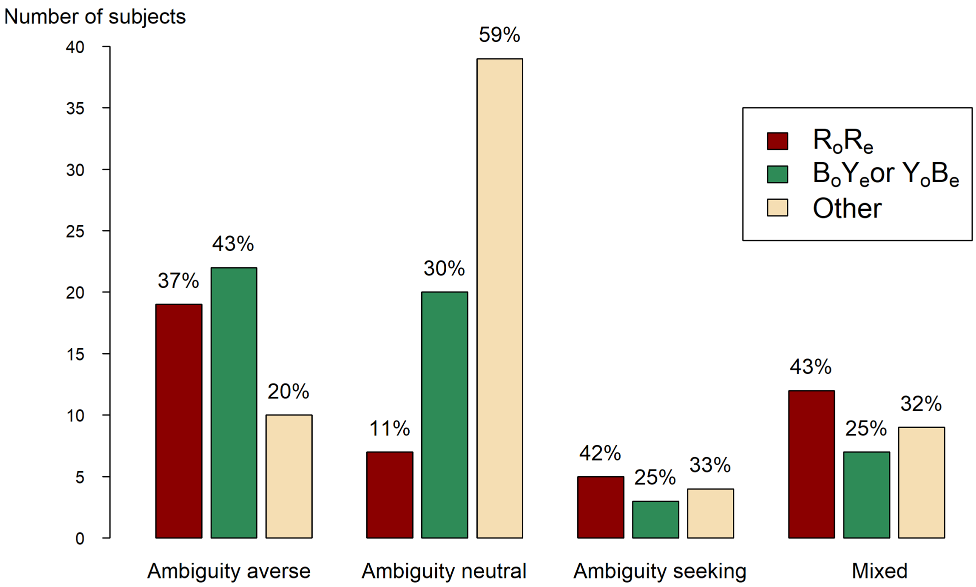


FIGURE 4. Results of decision situation 3

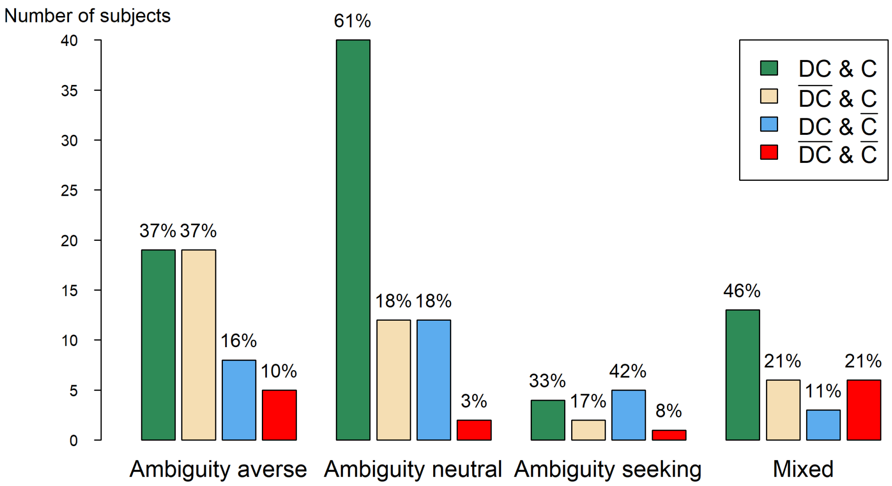


FIGURE 5. Subjects Satisfying Dynamic Consistency and Consequentialism Split Out by Ambiguity Attitude

respective property, while $\overline{DC}$ and $\overline{C}$ mean that the subject's choices violated the respective property.

One-hundred and four subjects, (66.2%), were dynamically consistent. Of the remaining fifty-three dynamically inconsistent subjects, twenty-six (49%) revised their initial choice of the *unconditionally* unambiguous option $B_oY_e$ or $Y_oB_e$ to the *conditionally* unambiguous option of choosing red.

One-hundred and fifteen subjects, (73.2%) satisfied consequentialism. In computing this number, we also included the seventeen subjects who switched between blue and yellow in decision situations 1 and 4 since we could not exclude the possibility that they were indifferent between betting on blue and beting on yellow.

Figure 5 also shows that ambiguity averse subjects were more prone to violate dynamic consistency than consequentialism (two-tailed proportion test, $p = 0.04$ , $N = 51$). Most ambiguity neutral subjects satisfied dynamic consistency and consequentialism and thus behaved accordng to subjective expected utility with Bayesian updating. The remaining ambiguity neutral subjects, who behaved according to subjective expected utility in decision situations 1 and 2, did not update in a Bayesian manner in decision situations 3 and 4. Among the ambiguity neutral subjects who violated at least one of the dynamic principles, there was no difference between the violation rates (two-tailed proportion test, $p = 1$, $N = 66$). The violation rates also did not differ among the ambiguity seeking subjects (two-tailed proportion test, $p = 0.40$, $N = 12$) or the mixed subjects (two-tailed proportion test, $p = 0.58$, $N = 28$). The proportions of ambiguity averse and ambiguity neutral subjects violating consequentialism were approximately equal (26% versus 21%). However, whereas 47% of the ambiguity averse subjects violated dynamic consistency, only 21% of the ambiguity neutral subjects violated dynamic consistency. Two-tailed Fisher Exact tests showed that the violations of dynamic consistency were associated with ambiguity attitudes ($p = 0.02$, $N = 157$), but the violations of consequentialism were not ($p = 0.19$, $N = 157$).

In light of Observation 1, it may seem surprising that nineteen ambiguity averse subjects satisfied both dynamic consistency and consequentialism. The key additional assumption underpinning that observation, however, is that the decision maker's preferences satisfy stochastic dominance. Indeed, eighteen out of the nineteen ambiguity averse subjects who satisfied both dynamic consistency and consequentialism violated stochastic dominance by choosing $R_oR_e$.[13]

Figure 6 presents the revision pattern of subjects who chose $B_oY_e$ or $Y_oB_e$ in decision situation 3. Among these subjects, only four (18%) ambiguity averse subjects satisfied dynamic consistency, whereas seventeen (77%) switched to the unambiguous option $R_o$ or $R_e$ in decision situation 4, in accordance with Observation 1. This pattern is similarly observed among the mixed subjects. However, the ambiguity neutral subjects exhibited different patterns. Only two (10%) switched to the unambiguous option. Half of the others satisfied dynamic consistency, whereas the other half switched from blue to yellow or vice versa.

## 5. Discussion

We have performed an experimental test of dynamic consistency and consequentialism, two key principles of dynamic decision making, and studied how these are related to ambiguity attitudes.

[13]The other subject opted for $R_oY_e$ in decision situation 3 and, when informed that the parity of the card drawn was odd, then selected red in decision situation 4. If the parity of the card drawn had been even, however, then in decision situation 4, opting for yellow or blue would have constituted a violation of consequentialism and opting for red a violation of dynamic consistency, again illustrating the point made in Observation 1 that in our study an ambiguity averse decision maker has to give up at least one of stochastic dominance, dynamic consistency, and consequentialism.

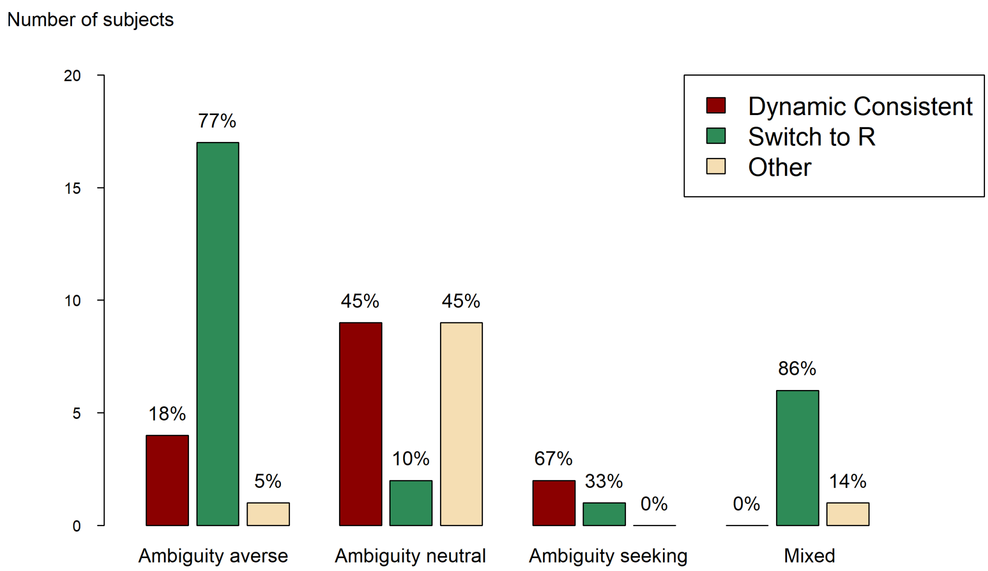


FIGURE 6. Revision Pattern of Subjects Choosing $B_oY_e$ or $Y_oB_e$ Split Out by Ambiguity Attitude

We used a variation of Ellsberg's three color problem in which the cards (used instead of balls) differed not only by color, but were also numbered. This allowed for a comparison between unconditional and conditional preferences, that is required to test consequentialism and dynamic consistency. We found that ambiguity averse subjects were nearly twice as likely to violate dynamic consistency as consequentialism. Sarin and Wakker (1998) showed that the only ambiguity model that can account for consequentialism and dynamic consistency simultaneously is maxmin expected utility. However, empirical tests provide little support for this model (e.g. Baillon and Bleichrodt 2015; Chew, Miao, and Zhong 2017). Hence, ambiguity models that strive for empirical realism should give up either dynamic consistency or consequentialism. Our results suggest that from a descriptive point of view, giving up dynamic consistency is the better choice.

Our test of consequentialism assumes that the decision makers view the two conditional decision situations and our no-history version of this decision situation as indistinguishable. We believe this extra assumption is plausible, but it is stronger than the definition in, for example, Ghirardato (2002) and, perhaps, our test may be interpreted as putting a lower bound on the support for consequentialism. In spite of this, we found that ambiguity averse decision makers were more likely to violate dynamic consistency than consequentialism.

In our tests, subjects did not know that after their initial choice they would receive a signal and would be asked to make another choice. One way to extend our research would be to make subjects aware of the sequential nature of the experiment. This would allow for a test of whether subjects are naive planners and use the same preference functional at each decision node, resolute planners, in which case they remain with their initial choice regardless of how uncertainty is resolved, or consistent planners, in which case they realize that they will not follow through some plans due to ambiguity aversion and, hence, delete these from their feasible sets.

Another way to extend our research is to use larger incentives. The cost for subjects of violating stochastic dominance, ambiguity neutrality, consequentialism and dynamic consistency was low, only 5 cents. If it would be driving subjects' choices then we would expect these to be random. However, we found no evidence of this providing some support for the conjecture that subjects took the questions seriously. It would be interesting to replicate our study with a larger penalty for ambiguity aversion and a larger cost of violating stochastic dominance. A limitation of using larger costs is that they may bias the results. If we put, for instance, a larger penalty on being ambiguity averse, then a substantial fraction of ambiguity averse subjects (who would like to choose $R$ in decision situation 1 of our experiment if there were no cost) may choose ambiguity neutral or seeking ($B$ or $Y$). This would deflate the support for ambiguity aversion. Hence, there appears to be a trade-off between the size of the penalty and the precision with which the degree of ambiguity aversion can be estimated. It is hard to say what the optimal balance is. Hopefully, future research can shed light on this question.

We found no relation between violations of consequentialism and ambiguity attitudes. It might be of interest to explore whether this also holds for settings other than decision under ambiguity. That is, do ambiguity neutral and ambiguity averse subjects who violate consequentialism also do so in decisions other than those made under ambiguity? This opens up a research agenda of the deeper causes of violations of consequentialism that we leave for future research.

Even though the comprehension test ensured that the subjects understood the combination of $B_oY_e$ and $Y_oB_e$ gives higher objective winning chance than $R_oR_e$, we nevertheless observed a non-negligible proportion of subjects preferring $R_oR_e$ in choice 3. This might be due to the relatively small cost of violating stochastic dominance. This also might be driven by an aversion to complexity. Options where the outcome depends on both the number and color of the card may appear more complex. In a similar situation, Dominiak and Schnedler (2011) report that subjects view a coin flip between two complementary ambiguous options as worse than either option on its own. This can be interpreted as evidence of complexity aversion in a related setting.

As we mentioned in the Introduction, we found the same proportion of subjects satisfying consequentialism as Dominiak, Duersch, and Lefort (2012) in spite of the differences in experimental design. We found more support for dynamic consistency than they did, but this was to a large extent driven by the subjects choosing the stochastically dominated option $R_oR_e$. As mentioned before, we believe these responses reflect more a cognitive shortcut to ambiguity aversion than an intrinsic preference for dynamic consistency. If we do not count these choices as reflecting dynamic consistency then the proportion of subjects satisfying dynamic consistency drops to 38.9%, which is close to the 32.2% observed by Dominiak, Duersch, and Lefort (2012). Moreover, among the ambiguity averse subjects 17.6% satisfy dynamic consistency, which, again, is very close to the 14.5% observed by Dominiak, Duersch, and Lefort (2012). The similarity between our findings and those of Dominiak, Duersch, and Lefort (2012) provides support for their robustness.

## 6. Conclusion

We have performed empirical tests of consequentialism and dynamic consistency, two key principles of rational dynamic choice. Most models of ambiguity aversion have to give up one of these two principles. We used a variation of the Ellsberg three urn decision problem to test these principles in the lab. Around 25% of our subjects behaved in line with subjective expected utility with Bayesian updating: they were ambiguity neutral and satisfied both dynamic consistency and consequentialism. Ambiguity averse subjects typically violated one of these principles, with violations of dynamic consistency about twice as common as violations of consequentialism .

## References

Baillon, A., and H. Bleichrodt (2015): "Testing ambiguity models through the measurement of probabilities for gains and losses," *American Economic Journal: Microeconomics*, 7(2), 77–100.

Chew, S. H., B. Miao, and S. Zhong (2017): "Partial ambiguity," *Econometrica*, 85(4), 1239–1260.

Cubitt, R., C. Starmer, and R. Sugden (1998): "Dynamic Choice and the Common Ratio Effect: An Experimental investigation," *Economic Journal*, 108, 1362–1380.

Dominiak, A., P. Duersch, and J.-P. Lefort (2012): "A dynamic Ellsberg urn experiment," *Games and Economic Behavior*, 75(2), 625–638.

Dominiak, A., and W. Schnedler (2011): "Attitudes towards Uncertainty and Randomization: An Experimental Study," *Economic Theory*, 48, 289–312.

Ellsberg, D. (1961): "Risk, ambiguity, and the Savage axioms," *The Quarterly Journal of Economics*, pp. 643–669.

Epstein, L. G., and M. Le Breton (1993): "Dynamically consistent beliefs must be Bayesian," *Journal of Economic Theory*, 61(1), 1–22.

Ergin, H., and F. Gul (2009): "A theory of subjective compound lotteries," *Journal of Economic Theory*, 144(3), 899–929.

Ghirardato, P. (2002a): "Revisiting Savage in a conditional world," *Economic Theory*, 20, 83–92.

——— (2002b): "Revisiting Savage in a conditional world," *Economic Theory*, 20(1), 83–92.

Gilboa, I., and D. Schmeidler (1989): "Maxmin expected utility with non-unique prior," *Journal of Mathematical Economics*, 18(2), 141–153.

——— (1993): "Updating ambiguous beliefs," *Journal of Economic Theory*, 59(1), 33–49.

Hammond, P. J. (1988): "Consequentialist foundations for expected utility," *Theory and Decision*, 25(1), 25–78.

Hanany, E., and P. Klibanoff (2007): "Updating preferences with multiple priors," *Theoretical Economics*, 2, 261–298.

——— (2009): "Updating ambiguity averse preferences," *The BE Journal of Theoretical Economics*, 9(1).

Johnson, C., A. Baillon, H. Bleichrodt, Z. Li, D. van Dolder, and P. P. Wakker (2020): "Prince: An Improved Method for Measuring Incentivized Preferences," forthcoming.

Karni, E., and Z. Safra (1990): "Behaviorally consistent optimal stopping rules," *Journal of Economic Theory*, 51(2), 391–402.

Kelsey, D., and S. LeRoux (2017): "Dragon Slaying with Ambiguity: Theory and Experiments," *Journal of Public Economic Theory*, 19, 178–197.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005): "A smooth model of decision making under ambiguity," *Econometrica*, 73(6), 1849–1892.

——— (2009): "Recursive smooth ambiguity preferences," *Journal of Economic Theory*, 144(3), 930–976.

Machina, M. (1989): "Dynamic consistency and non-expected utility models of choice under uncertainty," *Journal of Economic Literature*, 27, 1622–1668.

Machina, M. J., and D. Schmeidler (1992): "A More Robust Definition of Subjective Probability," *Econometrica*, p. 745–780.

Nau, R. F. (2006): "Uncertainty aversion with second-order utilities and probabilities," *Management Science*, 52(1), 136–145.

O'Donoghue, T., and M. Rabin (1999): "Doing it now or later," *American Economic Review*, 89(1), 103–124.

Sarin, R., and P. Wakker (1998): "Dynamic Choice and Non-Expected Utility," *Journal of Risk and Uncertainty*, 17, 87–120.

Schmeidler, D. (1989): "Subjective Probability and Expected Utility without Additivity," *Econometrica*, 57, 571–587.

Segal, U. (1990): "Two-stage Lotteries without the Reduction Axiom," *Econometrica*, 58(2), 349–377.

Seo, K. (2009): "Ambiguity and Second-Order Belief," *Econometrica*, 77(5), 1575–1605.

Siniscalchi, M. (2011): "Dynamic choice under ambiguity," *Theoretical Economics*, 6(3), 379–421.

Trautmann, S. T., and G. van de Kuilen (2015): "Ambiguity Attitudes," in *The Wiley Blackwell Handbook of Judgment and Decision Making*, ed. by G. Keren, and G. Wu, vol. 1, p. 89–116. Wiley-Blackwell.

## Appendix

This appendix presents the definition of consequentialism, which we test in our experiment.

Consequentialism is usually defined as an axiom about choices in decision trees. Hammond (1988) takes choices in decision trees as a primitive and derives preferences from these choices. An exception is Ghirardato (2002b), who defines consequentialism in terms of preferences. We shall present a definition of consequentialism as it applies to our experiments. For a definition in more general decision trees see Hammond (1988).

Consequentialism requires that individuals make the same choices in different trees with the same consequences. This applies to continuation trees, such as in our decision situation 4, as well as whole trees, such as in our decision situation 1.

Let $T$ denote a decision tree and let $C(T)$ denote the consequences which decision-maker would choose from tree $T$. Denote the decision tree starting from an odd (resp. even) node in the second experiment (i.e. decision situation 4 in Figure 2a), by $T_o$ (resp. $T_e$). Let $T_1$ denote the decision tree in Figure 2c, where implementer 1 puts $N$ blue balls into an urn which contains 33 red balls, and $67 - N$ yellow balls. Similarly, let $T_2$ denote the decision tree for the case where implementer 2 determines the composition of colors in the urn.

Consequentialism implies that

$$C(T_o) = C(T_e) = C(T_2).$$

We make the auxiliary hypothesis that

$$C(T_1) = C(T_2).$$

This says that subjects make the same choice when the urn is filled by implementer 1 as when it is filled by implementer 2. This can be justified by an appeal to symmetry of information.

We may thus derive

$$C(T_1) = C(T_o) = C(T_e). \tag{1}$$

This is our Observation 1.

If a decision-maker is ambiguity-averse (s)he will choose red in decision problem 1. If in addition we assume that the decision-maker is consequentialist then equation (1) implies that (s)he will also choose red in decision problem 4.

[Han Bleichrodt]Erasmus School of Economics, Rotterdam, The Netherlands and Research School of Economics, Australian National University, Australia.

[Jürgen Eichberger]Alfred Weber Institut, Universität Heidelberg, Germany.

[Simon Grant]Research School of Economics, Australian National University, Australia

[David Kelsey]Nottingham University Business School, England.

[Chen Li]Erasmus School of Economics, Rotterdam, The Netherlands.