Economics Department Discussion Papers Series

ISSN 1473 – 3307

A Test of Dynamic Consistency and Consequentialism in the Presence of Ambiguity

Han Bleichrodt
Jurgen Eichberger
Simon Grant
David Kelsey
Chen Li

Paper number 18/03

A TEST OF DYNAMIC CONSISTENCY AND CONSEQUENTIALISM IN THE PRESENCE OF AMBIGUITY

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

ABSTRACT. We test dynamic consistency and consequentialism, two key principles of dynamic decision making under ambiguity and relate violations of these principles to subjects' ambiguity attitudes. In our experiment, subjects received a signal which made it attractive for ambiguity averse subjects to deviate from their ex-ante contingent plan and violate dynamic consistency. We found that ambiguity averse subjects were indeed more likely to violate dynamic consistency than ambiguity neutral subjects, but not consequentialism.

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

JEL CLASSIFICATION. C72, D81

1. INTRODUCTION

In decision theory, ambiguity refers to decision problems in which the probability distribution over states of the world is itself unknown or uncertain. The term is due to Ellsberg (1961), who used it to describe "the nature of one's information concerning the relative likelihood of events." In the static models that have been developed to accommodate such ambiguity, the beliefs of a decision-maker either cannot be represented by a (single) probability measure or the probability measure is a second-order belief and the corresponding induced preferences

Date: September 13, 2018.

over compound lotteries do not satisfy the reduction of compound lotteries axiom.[1] In either case, the decision-maker's preferences are not *probabilistically sophisticated* in the sense of Machina and Schmeidler (1992).

To extend and apply these models to dynamic settings, however, raises questions not only about how such non-probabilistically sophisticated beliefs and the attendant (conditional) preferences should be updated but also, and arguably, more significantly, whether the updated preferences need only depend on outcomes on those states which are still possible (that is, exhibit *consequentialism*) and whether the associated behavior need be *dynamically consistent*.[2] Machina (1989) provides a careful and extensive discussion of the consistency issues arising in the context of dynamic decision making. The consistency issue is complicated by the fact that preference orders whose conditional preferences satisfy both properties (*consequentialism* and *dynamic consistency*) on an unrestricted domain are *probabilistically sophisticated* and do satisfy the *sure thing principle*.[3]

[1] 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 et al. (2005), Nau (2006) and Seo (2009), the subjective compound lottery model of Ergin and Gul (2009) as well as the experimental study of Halevy (2007).

[2] Formally, in the Savage framework, acts f are functions mapping states of a set S to a consequences in a set X. Unconditional preferences $\succcurlyeq$ and the family of conditional preferences $(\succcurlyeq_E)_{E \subseteq S}$ are defined on the set of all acts $\mathcal{F}$. In this model, *consequentialism* means that $f(s) = g(s)\,, s \in E$, implies $f \sim_E g$ for all $E \subseteq S$ and all $f\,, g \in \mathcal{F}$. *Dynamic consistency* means that, for $f(s) = g(s)\,, s \in S \setminus E$, $f \succcurlyeq g$ implies $f \succcurlyeq_E g$.

[3] For example, Epstein and Le Breton (1993) show that if the updated preferences satisfy consequentialism, then with an unrestricted domain of acts, dynamic consistency essentially entails the unconditional preferences are probabilistically sophisticated. Ghirardato (2002) shows that *consequentialism* and *dynamic consistency* imply subjective expected utility and Bayesian updating.

Given the impossibility of satisfying both principles for non-expected utility preferences, one strand of the literature imposes dynamic consistency as a property of the preferences themselves thereby precluding the need to consider how any potential dynamic inconsistency in the preferences might be manifested in the decision-maker's actual choices.[4] Another strand, which does allow for dynamically inconsistent preferences, models such individuals as sophisticated decision-makers who foresee such inconsistencies. So for them, dynamic consistency becomes an admissibility requirement for the plans of actions that the decision-maker considers *effectively* available to them to choose from. That is, they engage in what has been referred to in the literature as *consistent planning.*[5] An alternative for decision-makers with dynamically inconsistent preferences is to assume they do not foresee their dynamic inconsistency leading to sequential choices that turn out to be at odds with their initial plans.

Since a non-expected utility preference representation cannot satisfy *consequentialism* and *dynamic consistency* simultaneously, a natural question concerns the empirical relevance of these two properties. In this paper, we report the results of an experiment designed to test

[4] Hanany and Klibanoff (2007), Hanany and Klibanoff (2009), Klibanoff et al. (2009), Siniscalchi (2011) and Sarin and Wakker (1998) do this at the expense of dropping consequentialism. Epstein and Schneider (2003) show it is possible to retain both dynamic consistency and dynamic consequentialism if one restricts the domain of acts and conditioning events (or more precisely, filtrations) over which preferences are defined. Wang (2003) casts his analysis in a more complicated setting of consumption–information profiles, but effectively he is imposing similar restrictions to those of Epstein and Schneider (2003) on the domain of admissible problems.

[5] See for example Siniscalchi (2011).

whether sequential resolution of uncertainty in the presence of ambiguity is associated with consequentialist and (potentially) dynamically inconsistent preferences or dynamically consistent and (potentially) non-consequentialist preferences. The purpose of this experiment is not only to relate violations of consequentialism and dynamic consistency to the subjects' revealed ambiguity attitude but to gain insights into the reasons for such deviations. In our setup, we provide subjects with signals on which they can condition their actions and ask them for explicit choice of a conditional strategy. Once information is partially revealed by the signal, an ex-ante possibility to hedge against ambiguity may no longer exist. This changes the ambiguity experienced ex-ante from the ambiguity experienced after observing the signal. Hence, there is a rational for deviating from a conditional strategy for ambiguity-sensitive decision makers.

A novel feature of our set up is that conditioning on events is done via a signal (odd- or even-numbered balls) that allows us to use the strategy method to elicit the ex-ante conditional choice plans of subjects. These plans can in turn be compared with the subject's actual choices made after the signal's realization has been revealed. Another novel feature is that one of the available ex-ante ambiguity-free choice plans is (first-order stochastically) dominated as are all the ambiguity-free conditional choices available after the signal's realization. The introduction of the imposition of this small (probability) premium for selecting an ambiguity free option means we avoid the problem of potential indifference between ambiguous and ambiguous free options as encountered by other related studies such as Dominiak et al. (2012).[6] Furthermore, this allows us not only to detect violations of dynamic consistency and consequentialism but also violations of (stochastic) monotonicity.[7]

2. Ambiguity, information and dynamic inconsistency

Why might sequential resolution of uncertainty in the presence of ambiguity be natually associated with dynamically inconsistent choices by a consequentialist monotonic decision-maker? One answer to this question is provided by Eichberger et al. (2007) who, echoing sentiments originally expressed by Frisch and Baron (1988), note that ambiguity arises from uncertainty about probability created by missing information that is relevant and could be known. Hence, it seems plausible to assume that once our consequentialist decision-maker knows that an event has obtained, the only remaining ambiguity she faces relates to uncertainty about the *probabilities of subevents of that event*. Past (or borne) uncertainty that one may have had about the likelihoods of the counterfactual event and its subsets are no longer relevant. However, such uncertainty might have been relevant to the decision-maker at the time when she did not know whether or not the event had obtained. For instance, with partial resolution of uncertainty, an unambiguous event that was originally unambiguous may become conditionally ambiguous and vice versa. An ambiguity-averse consequentialist decision-maker, reacting to the change in ambiguity concerning an event, may thus exhibit dynamically inconsistent preferences.

[6] Hence our tests for non-neutral attitudes toward ambiguity are "robust" in the sense formalized by Grant et al. (2016).

[7] In this paper we shall refer to preferences as being monotonic if they agree with the partial ordering of (first-order) stochastic dominance.

To illustrate this point, consider a bag containing 200 cards numbered from 1 to 200. The cards numbered 1 to 66 are red. The next $2N$ cards (that is, those numbered from 67 to $66 + 2N$) are blue and the remaining cards (that is, those numbered from $67 + 2N$ to 200) are yellow. The only information the decision maker has about N is that it is a strictly positive integer not greater than 66.

The following 2×3 matrix summarizes the composition of the bag:

	RED	BLUE	YELLOW
ODD	33	N	$67 - N$
EVEN	33	N	$67 - N$

The decision maker has to pick a color to bet on, for odd cards and even cards. A card is to be drawn randomly from the bag. The decision-maker wins if the card matches one of the color-parity combination she chose. Let $[CC], C \in \{R, B, Y\}$ denote the bets placed by the decision maker, where the first and second C refers to the color bet for odd and even cards, respectively. For instance, if the decision maker bets on $[RB]$, she wins if it is either an odd red card *or* an even blue card. There are nine different possible combinations from which she may choose.

$$\begin{array}{ccc} [RR] & [RB] & [RY] \\ [BR] & [BB] & [BY] \\ [YR] & [YB] & [YY] \end{array}$$

Consider the pair of bets, $[RR]$ and $[BY]$. Given the information about the composition of the bag, the decision-maker knows she will win with $[RR]$ if any one of the 66 red cards in the bag are drawn. Similarly, she knows that she will win with $[BY]$ if any one of the $N + (67 - N) = 67$ cards that are either (odd and blue) or (even and yellow) are drawn. Thus if she views each card in the bag as equally likely to be drawn, then, since $67 > 66$, we should expect her to express a strict (unconditional) preference for $[BY]$ over $[RR]$.

Now suppose a card is drawn from the bag, but instead of simply being told which card it is, the decision-maker is only told the parity of the number on the card, that is, whether it is odd or even. What effect, if any, should this piece of information have on her conditional preferences over bets?

Let O (respectively, E) denote the event that the card drawn from the bag has an odd (respectively, even) number on it. If she were told the parity of the number of the card was odd, then recall that out of the one-hundred odd cards, she *knows* thirty-three are red, but she does not know the exact number of blue or yellow cards. So, if she were sufficiently ambiguity averse, then we might expect her to express a strict (conditional) preference for $[RY|O]$ over $[BY|O]$. And since she is a consequentialist which means she must be indifferent between the pair of conditonal bets $[RY|O]$ and $[RR|O]$, by transitivity we know this means she would strictly prefer $[RR|O]$ over $[BY|O]$, reversing her original unconditional preference.

By similar reasoning, if she were told the parity was even, we should expect her to express a strict (conditional) preference for $[BR|E]$ over $[BY|E]$ which again from her being a consequentialist entails a strict (conditional) preference for $[RR|E]$ over $[BY|E]$ also constituting a reversal of her original unconditional preference.[8]

[8] The strict (unconditional) preference for $[BY]$ over $[RR]$ combined with the pair of strict (conditional) preferences for $[RR|O]$ over $[BY|O]$ and $[RR|E]$ over $[BY|E]$ constitutes a violation of what Skiadas (1997) dubs *strict coherence.*

To summarize, before the card is drawn we expect every monotonic decision-maker to express a strict preference for $[BY]$ over $[RR]$. However, once the card has been drawn and the decision-maker has been informed of the parity of the card's number while remaining uncertain about its color, we would expect any consequentialist decision-maker sufficiently averse to betting on ambiguous odds *to reverse* her original preference.

Intuitively, the bet on $[BY]$ provides the decision maker with a natural hedge against her lack of knowledge concerning the color composition of the cards in the bag, thus resulting in her strict (unconditional) preference for $[BY]$ over $[RR]$. Being informed that the card drawn is odd (or even), however, destroys this hedge. Thus the originally favored bet now becomes ambiguous, whereas the originally dominated bet $[RR]$ remains unambiguous. Hence subjects will violate dynamic consistency, if they are sufficiently ambiguity-averse.

3. Experiment

We conducted a lab experiment to elicit subject's ambiguity attitudes and choices in the dynamic decision problem described in section 2.

Subjects. The experiment was computer-based, conducted in the ESE-Econlab at Erasmus University Rotterdam in March 2018. It consisted of 7 sessions, with 23 to 27 subjects participating in each session. In total 171 subjects were recruited from the ESE-Econlab subject pool. We collected preference data from 157 subjects, and the remaining 14 subjects (2 in each session) were randomly assigned to be implementers. The implementers assisted the experimenters to guarantee transparent and fair implementation of randomizations during the experiment (see section 3 for a description of implementers' tasks).

Stimuli. Subjects faced four decision situations during the experiment.

Decision situations $1a$ **and** $1b$**:** The subjects learned that implementer 1 wrote down a random number N between 1 and 67. Implementer 1 then put 100 cards into bag 1 with the color composition as shown in table 1.

	RED	BLUE	YELLOW
Frequency	33	N	$67 - N$

TABLE 1. Composition of Bag 1

In decision situation $1a$, subjects could choose to bet on red, blue or yellow. Betting on one color means that the subject would get 10 euros if at the end of the experiment, the color of the card drawn by implementer 1 matched the color that she bet on.

In decision situation $1b$, subjects could choose to bet on a combination of two colors, depending on her color choice in $1a$:

- If she chose to bet on red or blue in $1a$, then she could choose between betting on "red and yellow" and "blue and yellow";
- If she chose to bet on yellow in $1a$, then she could choose between betting on "red and blue" and "yellow and blue".

Betting on "red and yellow", for instance, in $1b$ means that the subject would get 10 euros if at the end of the experiment, the color of the card drawn by implementer 1 is either red or yellow.

Decision situations $2a$ and $2b$: The subjects learned that implementer 2 wrote down a random number M between 1 and 67. Implementer 2 then put 200 cards, numbered from 1 to 200, into bag 2 with the color and number composition as shown in table 2.

	RED	BLUE	YELLOW
Odd	33	M	$67 - M$
Even	33	M	$67 - M$

TABLE 2. Composition of Bag 2

In decision situation $2a$, subjects had to specify a color to bet on for both odd- and even-numbered cards. This means, a subject who, for instance bets on blue for odd and yellow for even, would win 10 euros if at the end of the experiment the number on the card drawn by implementer 2 was odd and its color was blue or its number was even and its color was yellow.

Before making their choices in $2a$, subjects had to answer three comprehension test questions before they could proceed. More information on these questions are included in the online appendix.

In decision situation $2b$, subjects were told that the number on the card drawn by implementer 2 was even (odd), depending on the card parity drawn by implementer 2. With the new information on the parity of number of the card drawn, subjects were asked to choose a color to bet on only for the even-(respectively, odd-)numbered cards that would matter for their payments.

Procedure and Incentives. The experiment was incentivized using the prior incentive system (Prince; Johnson et al. 2017). At the beginning of each session, every subject drew an envelop from a pile of n

sealed envelopes (n = number of participants in each session), on top of which a subject ID was written. Subjects who drew an ID starting with "m" became the implementers, and the others (with an ID starting with "p") participated in the main experiment.

Subjects in the main experiment learned that they would face different decision situations during the experiment. One of these decision situations was inside the envelope. They should not open their envelopes until the experimenters asked them to do so. At the end of the experiment, they would be paid for real according to their choices in the decision situation that came out of their envelopes.

The implementers followed one (of the two) experimenters to the implementer computers. They assisted the experimenters to construct the bags. The computer program recorded the random numbers they typed in, and the color and number of the card that they drew from bag 1 or bag 2. The parity of the card drawn from bag 2 was used to generate the proper decision situation $2b$ in the main experiment. At the end of the experiment, the implementers revealed the color and the number of the drawn cards to the other participants.

Implementers received a flat payment of 10 euros. Other participants received a show-up fee of 5 euros plus a variable amount depending on their choices. The average payment was 9.62 euros.

4. Results

Frequency of comprehension test failure. During the experiment, subjects had to pass a comprehension test to proceed with the experiment. Table 3 presents the frequency of the number of failures. In this section, we report the results with all subjects included. Removing subjects who fail the comprehension test multiple times does not affect our results (see online appendix).[9]

	0	1	2	3	4	5	7	8	9
frequency	55	47	22	15	10	2	3	2	1

TABLE 3. Frequency of comprehension test failure

Ambiguity attitude according to choices in $1a$ **and** $1b$**.** Table 4 presents the frequency of combinations of choices in decision situations $1a$ and $1b$, where each row (column) corresponds to a possible choice in $1a$ ($1b$).

	RB	RY	BY
R	0	23	46
B	0	5	34
Y	5	0	26

TABLE 4. Choice Pattern in 1a and 1b

Dominiak et al. (2012) classified subjects as ambiguity averse, ambiguity seeking, or ambiguity neutral according to their choices in the classical 3-color Ellsberg paradox. The decision situation that they presented to subjects were:

- $1a'$: betting on R or B.
- $1b'$: betting on RY or BY.

They found 54.4% ambiguity aversion (R and BY), 6.7% ambiguity loving (B and RY), and 38.9% ambiguity neutral (B and BY and R and RY).

Following the same classification logic, we deem subjects with choice combination R and BY as ambiguity averse (A); B and RY or Y and

[9] The online appendix is available at:
https://www.dropbox.com/s/2d5wr5fg7vb1y0i/onlineappendix.pdf?dl=0

RB as ambiguity seeking (S); B and BY or Y and BY as ambiguity neutral (N).

The combination of R and RY is not so straightforward to classify. As there is a small cost to choosing the unambiguous option R (33 versus $\frac{67}{2}$) the combination of choices R and RY cannot be classified as ambiguity neutral as in Dominiak et al. (2012). Since the choice of R in $1a$ suggests a subject is ambiguity averse but the choice of RY in $1b$ seems more akin to ambiguity seeking, we label this combination as reflecting "mixed" attitudes (M). One possible explanation for their choices, however, is that these subjects were ambiguity averse but did not realize that BY in $1b$ was a non-ambiguous option. We see support for this explanation below when analyzing the choices made by these subjects in decision settings $2a$ and $2b$.

Figure 1 plots the categorization of subjects' ambiguity attitudes according to their choice combinations in decision situations $1a$ and $1b$. Our finding is similar to Dominiak et al. (2012)'s finding, except that they observed more ambiguity aversion. The difference potentially comes from two sources. First, for the reason we noted above, subjects who we classified as "mixed" may actually be ambiguity averse. Second, in our setting, the ambiguous options were slightly more attractive than those in the standard Ellsberg (3-color) one-urn problem. In the standard Ellsberg one-urn problem, the proportion of Red was 1/3 while the proportion of Blue and yellow together was 2/3. In our setting, the divide was 0.33 and 0.67, respectively. Therefore, subjects had to be willing to incur a strictly positive "ambiguity (probability) premium" to choose the unambiguous options: R and BY.[10]

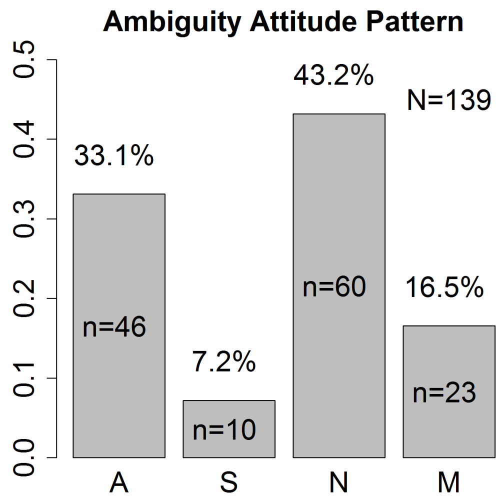

Figure 1. Ambiguity Attitude Pattern in decision situations 1a and 1b

	Frequency
RR	35
BY	33
YB	19
BB	21
YY	11
RB	8
BR	4
RY	7
YR	1

Table 5. Choice Pattern in 2a

Choice pattern in decision situation $2a$**.** Table 5 presents the frequency of choice combinations in situation $2a$. Figure 2 plots the categorization of subjects' choices in $2a$. The most striking feature is that

[10] Kelsey and LeRoux (2017) show that even small premia can have a big effect on the number of ambiguity-averse responses.

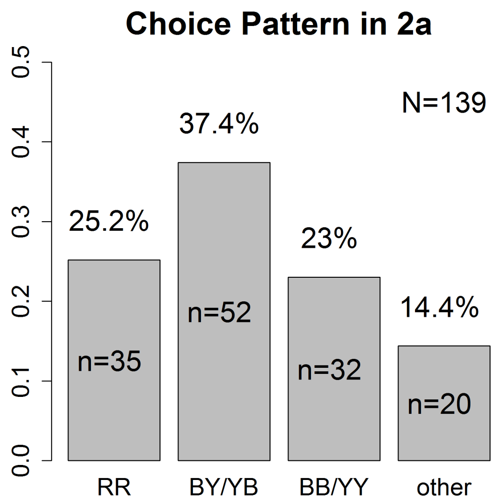

FIGURE 2. Choice Pattern in 2*a*

over a quarter of the subjects (27.4%) chose the combination RR, violating monotonicity. Recall, given a belief that any one card is just as likely to be drawn as any other, RR is strictly dominated by the combinations BY and YB. These two other unambiguous (but undominated) options were chosen by around a third of the subjects. A little over a fifth of the subjects (22.3%) revealed a belief that either blue or yellow was the most likely color, while the remaining 17.2% chose one of the other possible four combinations.

	RR	BY/YB	BB/YY	OTHER	TOTAL
A	16	22	2	6	46
S	4	3	1	2	10
N	6	20	27	7	60
M	9	7	2	5	23
Total	35	52	32	20	139

TABLE 6. Choice Pattern in 2a by Ambiguity Attitude

Table 6 breaks down the choice pattern of $2a$ by the ambiguity attitude exhibited by the subjects' selections in choice situations $1a$ and $1b$. Looking across the first row, we see that 80.4% of the ambiguity averse subjects choose one of the three unambiguous options (RR, BY and YB), more than the proportion (40.9%) among ambiguity neutral subjects (p-value < 0.01 in proportion test). However, it is not significantly different to the proportion among subjects classified as "mixed" (p-value $= 0.33$). We do not include ambiguity seeking subjects in the comparison as the power of statistical tests are low due to the low number of ambiguity seeking subjects.

Looking down the first column we see that 44.1% of those choosing the unambiguous (but dominated) option RR had been classified as ambiguity averse and this proportion increases to 72.1% if we include with the ambiguity averse those who had been classified as "mixed" from their choices in $1a$ and $1b$. Recall, we conjectured that some of those classified as "mixed" might actually be ambiguity averse but failed to realize that BY in choice situation $1b$ was an unambiguous option. Consistent with this hypothesis, we see from looking across the fourth row that 42% of those classified as "mixed" chose the dominated and most easily recognized unambiguous option RR while over two-thirds of them chose one of the three unambiguous options.

In summary, the classification of ambiguity attitude from choices made in situations $1a$ and $1b$ accords well with our expectation for the pattern of the subjects' choices in situation $2a$. Ambiguity averse individuals predominantly selected one of the three unambiguous options, albeit with a sizeable fraction selecting the dominated RR. Ambiguity neutral subjects predominantly choose either one of the undominated options or "to bet" either on blue or on yellow. Only a small fraction selected the dominated option RR. Of those who selected the dominated option RR, almost one-half had been classified as ambiguity averse and a further sizeable fraction had been classified as "mixed" with this choice of RR, adding weight to our conjecture that these subjects were also ambiguity averse but had failed to realize BY was a non-ambiguous option for choice situation 1b.

Dynamic Consistency and Consequentialism. Turning to choice situation 2b, notice that the subjects' choices in 2a and 2b provide a test of dynamic consistency while their choices in 1a and 2b provide a test of consequentialism.

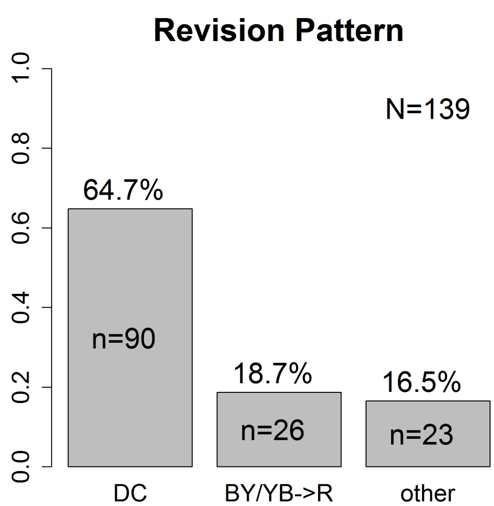

FIGURE 3. Choice Revision from 2a to 2b

Figure 3 plots the categorization of subjects' revision behavior from 2a to 2b. Overall, nearly two-thirds (66.2%) were dynamically consistent. Among those who were dynamically inconsistent, just under one half (26/53) revised their initial choice of the unambiguous option BY or YB to the *conditionally* unambiguous option R.

	R	B	Y
R	49	12	8
B	9	23	7
Y	8	8	15

TABLE 7. Choice Pattern in 1a and 2b

Table 7 presents the frequency of combinations of choices in situations 1a and 2b, where each row (respectively, column) corresponds to a possible choice in 1a (respectively, 2b). Choice patterns lying on the diagonal in table 7 (that is, betting on the same color in both situations) are deemed consistent with consequentialism. We also classify the choices of the seventeen subjects who chose blue in one choice and yellow in the other as also consistent with consequentialism, since their choice in 1a does not exclude the possibility that they viewed the drawing of a blue card as being as equally as likely as the drawing of a yellow card. Hence, in sum, nearly three-quarters (115/157 or 73.2%) of the subjects made choices consistent with consequentialism.

	A	S	N	M	Total
DC&C	16	3	37	9	65
n_DC&C	18	1	12	6	37
DC&n_C	8	5	9	3	25
n_DC&n_C	4	1	2	5	12
Total	46	10	60	23	139

TABLE 8. Dynamic Consistency and Consequentialism

Table 8 presents the overall frequency of subjects satisfying dynamic consistency and/or consequentialism for each of the four catergories of ambiguity attitude. Overall as can be gleaned from the last column in table 8), 66.2% were dynamically consistent while 73.2% satisfied consequentialism. Subjects were not more likely to violate dynamic consistency than consequentialism (p-value = 0.22 in McNemar test). Comparing the second and the third row of the table, we see that ambiguity averse and mixed attitude subjects were more prone to violate dynamic consistency than consequentialism while ambiguity neutral subjects were equally prone to violate dynamic consistency and consequentialism. Overall, ambiguity neutral subjects were the most likely to satisfy both dynamic consistency and consequentialism, and at the same time the least likely to violate both principles. Nevertheless, 39% (26 out of 66) seemingly "rational" ambiguity neutral subject in the static decision situation violated at least one of the rationality principles in the dynamic setting, suggesting that they are not true Bayesians.

At first sight, it may seem surprising that nineteen ambiguity averse subjects satisfied both dynamic consistency and consequentialism, as the literature suggests that a decision maker cannot satisfy dynamic consistency and consequentialism while being ambiguity averse. The hidden assumption, however, is that the decision maker's preferences satisfy (stochastic) monotonicity. Previous tests of dynamic decision principles focused only on dynamic consistency and consequentialism. Our design is the first to allow for joint detection of violations of dynamic consistency, consequentialism and (stochastic) monotonicity. Eighteen out of the nineteen ambiguity averse subjects who satisfied both dynamic consistency and consequentialism chose the dominated option $[RR]$ in situation 2a, violating monotonicity.[11] For these subjects, ambiguity aversion led to a direct cost on their payoffs. We clearly demonstrate that, in the dynamic setting, something has to give for being ambiguity averse: stochastic monotonicity, dynamic consistency, or consequentialism.

	A	S	N	M
DC	0.48	0.20	0.23	0.48
C	0.26	0.60	0.18	0.35
Mon	0.35	0.40	0.10	0.39

TABLE 9. Violation of Decision Principles by Ambiguity Attitude

Table 9 presents the violation rate of each decision principle. Consistent with our conjecture, ambiguity averse and "mixed" subjects are as equally likely to violate dynamic consistency (p-value = 0.90 in proportion test), consequentialism (p-value = 0.71 in proportion test) and monotonicity (p-value = 0.81 in proportion test). However, compared to ambiguity neutral subjects, ambiguity averse subjects are more prone to violate dynamic consistency and monotonicity (p-value < 0.01 in proportion tests) but not so for consequentialism (p-value = 0.75 in proportion test).

Next we take a closer look at the relationship between the subjects' initial choices in $2a$ and whether they satisfied dynamic consistency or consequentialism.

Table 10 presents dynamic consistency by initial choice in $2a$. Among the dynamically consistent subjects, the most common initial choices are RR, BB and YY. Whereas, among the dynamically inconsistent

[11] The one subject who did not fail our monotonicity test chose '$[RY]$ in 2a and 'R' in 2b for an odd card. Had the card drawn been even, this subject would have violated either consequentialism or dynamic consistency. Our test, in this sense underestimates the proportion of violations.

subjects, nearly seventy per cent (37/53) chose BY or YB. And as we see from Table 11, seventy per cent of these (26/37) revised their choice to the conditionally unambiguous option R. Furthermore, the overwhelming majority of these (23/26 or 88.5%) were A or M. This finding supports our theoretical prediction that when the partial resolution of uncertainty alters the ambiguity of an event, ambiguity averse people will react to this change and exhibit dynamic inconsistency.

	DC	n_DC
RR	31	4
BY/YB	15	37
BB/YY	29	3
other	15	5

TABLE 10. Dynamic Consistency by choice pattern in 2a

	BY/YB->R	rest
A	17	29
S	1	9
N	2	58
M	6	17

TABLE 11. Revision pattern by ambiguity attitude

5. Conclusion

This paper presents an experimental study of how non-expected utility preferences are revised when new information is received. Our experiment uses a novel information structure, parity of numbers as a signal, in order to elicit preferences over signal-dependent lottery choices. These revealed preferences can be compared with the choices made once the subjects had learned the signal.

Behavior in dynamic choice is related to ambiguity attitude as measured in our experiment. Compared with ambiguity neutral subjects, ambiguity averse ones were more likely to violate dynamic consistency but not consequentialism. A large proportion of ambiguity averse subjects who violated dynamic consistency were due to the change in ambiguity before and after the signal revelation. Subjects classified with mixed ambiguity attitudes behaved similarly to ambiguity-averse subjects. As might be expected, ambiguity-neutral subjects were least likely to violate either consequentialism or dynamic consistency. Thus their behavior was the closest to subjective expected utility theory. There were not enough ambiguity-loving subjects to draw clear conclusions.

Modeling sequential decisions of agents with non-expected utility preferences in economics, there are essentially three responses to the unavoidable conflict between dynamic consistency and consequentialism.

Naive choice. The individual uses the same preference functional at each decision-node. In this case, the best option is chosen disregarding links to past or future decisions. Naive choice satisfies consequentialism but fails to be dynamically consistent. Naive choice would predict ambiguity-averse subjects choosing an undominated unambiguous plan of action BY or YB in 2a but switching to R in 2b. This was indeed the most common pattern of switching observed in our experiment.

Resolute choice. Individuals formulate an optimal strategy at the initial node and remain with it regardless of how uncertainty is resolved. Resolute choice satisfies dynamic consistency but violates consequentialism. In our experiment an ambiguity-averse individual practicing resolute choice would choose one of the two dominant and unambiguous options BY or YB and would not revise his/her choice in 2b.

Consistent Planning. Consider an ambiguity-averse individual satisfying consequentialism. Knowing that she will choose R over B or Y if faced with these options, she realizes that a plan BY is not sustainable without a commitment device. A consistent planner would take this into account, discard the contingent plans BY and YB from his set of feasible options, and choose the preferred plan from the remaining alternatives. An effort to choose consistent plans may explain why some ambiguity-averse individuals chose the dominated action RR in 2a. Note, however, that our experimental setup does not provide a definitive test for this hypothesis since subjects were not told in advance which questions they would be asked in 2b.

Our observations suggest that subjects seem to follow neither naive nor resolute choice. This may be viewed as indirect evidence in favor of consistent planning. Thus, our results suggest a direct test of individual efforts to make consistent consequentialist choices as a topic for further research.

References

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

Eichberger, J., S. Grant, and D. Kelsey (2007). Updating Choquet Beliefs. *Journal of Mathematical Economics 43*(7-8), 888–899.

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

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

Epstein, L. G. and M. Schneider (2003). Recursive multiple-priors. *Journal of Economic Theory 113*(1), 1–31.

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

Frisch, D. and J. Baron (1988). Ambiguity and rationality. *Journal of Behavioral Decision Making 1*(3), 149–157.

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

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

Grant, S., J. Kline, I. Meneghel, J. Quiggin, and R. Tourky (2016). A theory of robust experiments for choice under uncertainty. *Journal of Economic Theory 165*, 124–151.

Halevy, Y. (2007). Ellsberg Revisited: An Experimental Study. *Econometrica 75*(2), 503–536.

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

Hanany, E. and P. Klibanoff (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 (2017). Prince: An Improved Method for Measuring Incentivized Preferences.

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.

Klibanoff, P., M. Marinacci, and S. Mukerji (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*, 745–780.

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

Sarin, R. and P. P. Wakker (1998). Dynamic choice and nonexpected utility. *Journal of Risk and Uncertainty 17*(2), 87–120.

Schmeidler, D. (1989, May). 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.

Skiadas, C. (1997). Conditioning and aggregation of preferences. *Econometrica* (2), 347–367.

Wang, T. (2003). Conditional preferences and updating. *Journal of Economic Theory 108*(2), 286–321.

Affiliations

(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) Australian National University and The University of Queensland, Australia

(David Kelsey) Department of Economics, University of Exeter, England.

(Chen Li) Erasmus School of Economics, Rotterdam, The Netherlands.