

TWO OUT OF THREE AIN'T BAD: A COMMENT ON “THE AMBIGUITY AVERSION LITERATURE: A CRITICAL ASSESSMENT”

MARCIANO SINISCALCHI

Northwestern University

Al-Najjar and Weinstein (2009) propose to scrutinize the implications of recent theories of ambiguity in dynamic settings. They conclude that such implications are so unreasonable as to cast doubts on the legitimacy of the theories under consideration. The present paper argues that the seemingly unreasonable implications highlighted by Al-Najjar and Weinstein can be understood as the result of basic trade-offs that arise naturally in the presence of ambiguity. In particular, Al-Najjar and Weinstein are uncomfortable with the possibility that an ambiguity-averse individual may reject freely available information; however, this phenomenon simply reflects a trade-off between the *intrinsic* value of information, which is positive even in the presence of ambiguity, and the value of *commitment*.

1. INTRODUCTION

Nabil Al-Najjar and Jonathan Weinstein (Al-Najjar and Weinstein, 2009, NW henceforth) offer a thoughtful critical assessment of the literature on decision under ambiguity originating from the seminal contributions of Schmeidler (1989) and Gilboa and Schmeidler (1989). This literature seeks to generalize Leonard Savage’s axiomatization of subjective expected utility (Savage, 1954) so as to accommodate phenomena such as the Ellsberg paradox (Ellsberg, 1961). Following Savage’s lead, it adopts the *subjective* point of view: the decision-maker’s preferences are the

This title is borrowed from a song by Jim Steinman, recorded by the band Meat Loaf in 1977.

fundamental input to the analysis. Consequently, specific functional representations, such as Savage's expected utility (EU) or Gilboa and Schmeidler's maxmin-expected utility (MEU), are merely convenient mathematical stand-ins for the underlying preferences. Furthermore, the individual's attitudes toward ambiguity are viewed as a matter of taste, just like her attitudes toward risk.

The subjective ambiguity literature has another feature in common with Savage's axiomatization of SEU. Most theories of choice under ambiguity are initially formulated and axiomatized in a *static* decision setting: the decision-maker (DM) takes an action, then uncertainty is resolved and the DM's payoff is determined. Extensions to *dynamic* choice are typically provided in subsequent contributions.

NW propose that it is important to flesh out the implications of subjective, static theories of ambiguity in *dynamic* decision problems. This strikes me as an important objective: almost all potential applications of interest in economics involve some dynamic element. Furthermore, static expected-utility theory comes equipped with a natural, essentially "built-in" theory of updating and dynamic choice; it is quite natural to ask whether existing theories of ambiguity also allow a similarly convenient and effective analysis of dynamic behaviour.

The main point that NW make is that subjective theories of ambiguity lead to problematic behaviour in dynamic-choice problems – so problematic, in fact, that even a DM who is disposed to exhibit the modal preferences in the Ellsberg paradox should actually *reconsider them*, once she takes into account their distasteful implications for dynamic choice:

the scrutiny of dynamic settings . . . reveals the extent to which a decision maker ought to view the Ellsberg choices as absurd and embarrassing.

My response, in a nutshell, is this: perhaps NW themselves would be embarrassed by these choices, but there is no fundamental canon of rationality according to which *every* DM should feel similarly uncomfortable. Each pattern of behaviour documented by NW (all of which, I should add, are well-known in the literature) is perfectly reasonable if interpreted within a specific framework for dynamic choice under ambiguity. Thus, NW's "embarrassment" is an expression of their own, subjective (meta)preferences over alternative modelling choices. Other decision-makers and modellers need not feel similarly ashamed. I, for one, do not.

I warn the reader that, in order to make the arguments precise, I will need to delve a bit deeper into the structure and formalities of dynamic decision-making than NW do. Once the mechanics of choice behaviour are represented with sufficient precision, all important issues and trade-offs emerge clearly. I view this as a necessity given the subject matter at hand: informal commentaries on the behaviour predicted by theories of

ambiguity-sensitive choice can easily overlook (or fail to emphasize, or perhaps conceal?) subtle aspects that are crucial to understanding the motivation behind that behaviour. Readers interested in a full-blown, formal development of the ideas I employ here may wish to consult Siniscalchi (2009).

This paper is organized as follows. Section 2 contains some preliminaries: it reviews the essential properties of dynamic choice under EU, argues that these *must* be relaxed to accommodate ambiguity, and summarizes the different ways they have been relaxed in the literature. Also, Section 2.4 revisits the “rationality test” adopted by NW, based on the criterion that the DM should not be “ashamed” of her choices. I fully subscribe to this criterion, and indeed argue that one should push this approach even further than NW do.

I then turn to the main substantive counterarguments vis-à-vis NW. Section 3 discusses *sophisticated choice*, with special emphasis on its implications for the *value of information* under ambiguity. I hope I will clarify that the alleged “information aversion” seemingly exhibited by ambiguity-sensitive, sophisticated individuals is, in actuality, the result of a rational trade-off between the value of information and the *value of commitment*.

I then offer some comments on *non-consequentialist choice* in Section 4. While I share some of NW’s misgivings as regards relaxing consequentialism, I do not find their arguments especially compelling. I discuss their “sunk-cost” example, and propose an alternative one that, in my view, makes a rather stark point about the unpleasant implications of relaxing consequentialism precisely when ambiguity is an issue.

2. PRELIMINARIES, AND WHY TWO OUT OF THREE AIN’T BAD

2.1 Terminology and notation

I will attempt to adhere to the notation and terminology in NW. In particular, I assume that the definition of *states*, *consequences* and (Savage) *acts* are understood. Prior and conditional preferences given an event E are denoted by \succsim and \succsim_E respectively.

Throughout this paper, dynamic-choice problems will be represented more concretely by means of decision trees. I shall refer to *nodes*, *actions* available at a node, and *terminal nodes*; these terms all have standard meaning, but an example may still help. Following NW, we consider a dynamic version of the three-colour-urn Ellsberg example. The underlying urn contains 30 black balls and 60 red or yellow balls, in unspecified proportions. If the DM chooses L , then he gets to choose again between u and d without further information. If, however, the DM chooses R , she gets 10 in case the ball drawn is yellow; otherwise, she learns that the ball drawn is black or red, and again gets to choose between u and d .




FIGURE 1. A decision tree (NW fig. 5).

Figure 1 corresponds to NW's Figure 5; I eliminated "information sets" by depicting the DM's initial choice before the (partial) realization of the uncertainty. Filled circles (\bullet) denote decision nodes, and empty ones (\circ) indicate chance nodes, where uncertainty is (partially) resolved. The initial node is represented by a "fatter" filled circle. The actions available at the initial node are L and R ; following L , the DM chooses again between the actions u and d , after which all uncertainty is resolved. Following R , uncertainty is partially resolved: the DM learns whether the ball drawn is yellow or not. If it is, she obtains a payoff of 10; otherwise, she gets to choose between u and d .

To describe the DM's *intended* choices throughout the tree, we consider the notion of a *plan of action*, or simply "plan": a specification of actions at all decision nodes that are not ruled out by prior actions.¹ In Figure 1, the possible plans of action are (L, u) , (L, d) , (R, u) and (R, d) .

A plan of action is, strictly speaking, not a Savage act; we can, however, apply the principle of *reduction* to "map" it to a Savage act. Specifically, for each possible state, we trace out the path through the tree generated by the plan, until a terminal node is reached and the resulting consequence is determined. For instance, the plan (L, u) corresponds to the act $(10, 0, 10)$, where (x, y, z) is the act that delivers prize x in state b , prize y in state r , and prize z in state y . Similarly, the plan (R, d) maps to the act $(0, 10, 10)$.² In the tree of Figure 1, the actions u and d at the non-initial decision nodes also map to acts in a natural way. Specifically, the action u at the decision node

¹ If some action at a node prevents a subsequent node from being reached, regardless of the realization of the uncertainty, then the DM's counterfactual choices at that unreached node cannot affect preferences. Thus, it is appropriate to focus on plans of actions.

² Reduction is a substantive assumption on preferences: see e.g. Kreps and Porteus (1978). We shall adopt it to conform to NW, and also in keeping with most applications of EU to dynamic problems.

following L maps to the act $(10, 0, 10)$; the action u at the node following R and the event $\{b, r\}$ instead maps to a partial act $(10, 0, *)$, where we do not specify the outcome in case the ball drawn is yellow because, at that node, the ball is *known* to be either black or red.

2.2 Dynamic choice with EU preferences

As noted above, static EU preferences come equipped with a “built-in” theory of dynamic choice. Among its many virtues, Savage’s sure thing principle provides the behavioural underpinnings for the usual definition of conditional probabilities (i.e. Bayesian updating). To elaborate, given a (non-null) event E , stipulate that an act f is preferred to another act g given E , written $f \succ_E g$, if and only if $fEh \succ gEh$ (in the usual notation for composite acts). For this definition of conditional preferences to be well-posed, it must be the case that we cannot find two acts h, k such that $fEh \succ gEh$ and $fEk < gEk$: the sure thing principle is precisely the statement that two such acts cannot be found. In other words, the sure thing principle can be restated as follows: the above definition of conditional preferences (henceforth “Savage’s rule”) is well-posed.³ Finally, it is immediate to verify that, if P is the probability that represents prior preferences jointly with a utility function u , then $f \succ_E g$ if and only if $E_P[u \circ f | E] \geq E_P[u \circ g | E]$.

Complementing this construction of conditional preferences with the reduction assumption, we thus obtain a complete theory of behaviour in dynamic decision problems. Furthermore, dynamic EU behaviour enjoys two key properties.

- *Consequentialism*: preferences at a given decision node in the tree are fully determined by the residual uncertainty and the payoffs that may be obtained starting from the node under consideration. In other words, subtrees can be analysed independently of the whole tree; foregone payoffs, unrealized events, and actions taken in the past, as well as intended future actions, do not influence conditional preferences.
- *Dynamic Consistency*: at any decision node, the DM will be willing to carry out the plan of action that she determined to be optimal *ex ante*. In other words, if the DM deems a given action available at a node to be optimal when she considers the entire decision tree, she continues to prefer it to all other available actions upon actually reaching that node.

These properties are both attractive and convenient: for instance, they provide the basis for backward induction and dynamic programming, two techniques whose practical importance simply cannot be overstated.

³ Indeed, this precisely how Savage himself restates this property in the back cover of the Dover edition of the *Foundations of Statistics*.

Both properties descend from Savage's updating rule and the sure thing principle. In fact, it turns out that Consequentialism and Dynamic Consistency are *equivalent* to Savage's rule and the sure thing principle: see e.g. Ghirardato (2002), or Siniscalchi (2009, §4).

2.3 Ambiguity and dynamic choice: three approaches

This last observation is the crux of the matter. As is well-known, ambiguity, as manifested for instance in the Ellsberg paradox, entails precisely a violation of the sure thing principle. It then follows that, *if preferences are ambiguity-sensitive, then either Consequentialism or Dynamic Consistency (or both) must fail in some respect*. A recent, illuminating discussion is provided by Hanany and Klibanoff (2007b); see also Ghirardato (2002) and Siniscalchi (2009). Furthermore, it is well-known that a similar statement holds for non-expected utility models of choice under risk: see e.g. Hammond (1988, 1989), Machina (1989), Karni and Safra (1989), and Karni and Schmeidler (1991). The underlying idea is similar: Consequentialism and Dynamic Consistency are *separability* properties, and so are the sure thing Principle (in the setting of choice under uncertainty) and the Independence axiom (for choice under risk); thus, one should expect a tight connection between them (under the reduction assumption: see Segal, 1990).

To sum up, the modeller interested in ambiguity can essentially pick *any two out of three ingredients*: full generality in the representation of ambiguity attitudes, Consequentialism, or Dynamic Consistency. She *cannot* have all three (however, she can have a little bit of each). This motivates the title of this paper. It is important to reiterate that this "impossibility result" is, essentially, a folk theorem, as the above references indicate; versions of the examples presented by NW are also well-known.

The subjective ambiguity literature has explored several combinations of these three ingredients. The approach pioneered by Epstein and Schneider (2001) has proved most popular in applications. It maintains Consequentialism and *restricts*, but does not drop, Dynamic Consistency. By necessity, this approach also restricts the range of attitudes toward ambiguity that may be represented. I do not discuss this approach in this paper, but NW do so in §3.4. Hanany and Klibanoff (2007b) advocate dropping Consequentialism; their approach is discussed by NW in §3.3, and here in Section 4. The advantage is that one maintains the full power of Dynamic Consistency, and full generality in the representation of ambiguity attitudes; the drawback is that conditional preferences must depend upon the entire "context" of the decision, i.e. the full tree and the intended plan of action.

Yet another approach was advocated by Siniscalchi (2009) in the present context, but is inspired by earlier work in the context of certain

and risky choice by Strotz (1955–6), Gul and Pesendorfer (2005) and Karni and Safra (1989) among others. It retains Consequentialism and drops Dynamic Consistency in its entirety, thereby allowing for full generality in the representation of ambiguity attitudes. To resolve conflicts among preferences at different decision points, this approach adopts the principle of *sophistication* (or more precisely Consistent Planning). I devote Section 2.4 to this approach, because I find the discussion in NW (§3.2) lacking in detail and, as a result, ultimately misleading.⁴

The details of these approaches will be considered below. However, it is important to emphasize that, in each case, one or more properties of dynamic EU behavior are *consciously and purposely* relaxed, and a coherent theory of dynamic choice is assembled out of the remaining properties. By “coherent” I mean a theory wherein every preference can be rationalized in an internally consistent way, invoking the properties that are maintained and pointing to those that are relaxed. For instance, NW complain that a DM who follows the approach advocated by Hanany and Klibanoff (2007b) ends up “distorting” her beliefs depending on her prior actions, so as to maintain Dynamic Consistency; they assert that “[t]here is no independent motivation for why a rational decision maker would ever engage in such distortions”. But allowing conditional preferences to depend upon past actions (and the intended continuation plan) is the very *nature* of the non-consequentialist approach advocated by Hanany and Klibanoff. Thus, for a DM who follows their approach, it is not at all “irrational” to engage in such “distortions”: it is precisely what she is supposed to do!

Of course, it is perfectly legitimate for NW to express *their own subjective distaste* for such distortions. However, seen in this light, NW’s “critique” is revealed for what it is – a statement about their own modelling preferences. This turns out to be true for almost all of their critical observations.

2.4 Introspection-proof preferences

Before tackling sophistication and non-consequentialist choice, I offer a general observation on the basic idea of “rationality” in the present context.

As NW note, there is a bit of circularity in discussing the rationality of preferences in the context of an axiomatic decision model (or, equivalently, one that can be provided with axiomatic foundations): the *axioms* themselves are meant to tell us what is “rational behaviour”.

Here, I differ only slightly from NW. In my view, for a preference relation to be deemed “rational”, it must be at least transitive and monotonic (i.e not leading to strictly dominated choices). For this reason,

⁴ NW also consider (§3.1) another approach, which assumes *naiveté* instead of sophistication. I view naiveté as a form of bounded rationality; for instance, *unlike sophistication*, it can lead to strictly dominated choices. For this reason, I do not discuss this approach here, and fully agree with NW’s critical assessment of it.

I concur with NW's assessment of naiveté (§3.1): it cannot be considered a rational theory of dynamic behaviour.

However, arguing about other axioms is quite a bit more delicate. To circumvent this difficulty, NW propose that preferences may be deemed "rational" if they are *immune to introspection*. That is, upon reconsidering the preferences she expresses in a given experimental or actual decision setting, the DM is not ashamed of them, and will not want to revise them.

I find this approach eminently reasonable, and would actually push it even further than NW do. For the purposes of understanding the issues at hand, the decision theorist should *only* take into account preferences that are immune to introspection. In other words, we should first give the DM an opportunity to "rethink her preferences", and proceed with the analysis only after she has done so.

This is not just a minor quibble. Suppose, for instance, that even after careful reconsideration, the DM's prior and conditional preferences are dynamically inconsistent. In particular, the DM *may* have constructed her preferences using some model of choice under ambiguity, augmented with an updating rule; if so, we assume that the DM has taken the time to "see how she feels about them", and ultimately decided that those are, in fact, her preferences. If this is the case, then we as decision theorists simply have to accept the DM's dynamic inconsistency. We can, of course, express our own *subjective judgment* as to whether *we* would feel ashamed of such inconsistency, even if the DM does not. But, in doing so, we have to be very clear on one fact: we are expressing *our own preferences*. Otherwise, we would not be identifying "rationality" with "introspection-proofness": we would be imposing our own canons of reasonable behaviour onto the DM.

There is a further corollary to this methodological stance. If we are trying to criticize the way this DM copes with her dynamic inconsistency, our arguments *cannot* boil down to lamenting the fact that the DM is dynamically inconsistent! I will refer back to this observation several times in the following.

3. DYNAMIC INCONSISTENCY AND SOPHISTICATION

3.1 NW's example: a more precise analysis

I shall now focus on NW's discussion of sophisticated choice and "information aversion". The starting point of their analysis is the observation that (as is well-known) if prior and conditional preferences are dynamically inconsistent, one needs to stipulate how the DM resolves this conflict. NW quickly discard the naive-choice approach, and I have nothing to add to their critique. They then turn to *sophisticated choice*. In a two-period decision problem, a sophisticated DM holds correct (prior) "beliefs" about her own time-2 choices, and takes them into account when

choosing an action at time 1. NW mainly discuss the example reproduced in Fig. 1 above (Fig. 5 in their paper), and I will for the most part stick to their example. I also discuss richer and more insightful examples in §3.3 below.

NW consider a DM with *ex-ante* MEU preferences, with a set of priors C consisting of all probabilities P on the state space $\{b, r, y\}$ with $P(b) = \frac{1}{3}$, and prior-by-prior updating. Observe that, for all $P \in C$, $P(b|\{b, r\}) \in [\frac{1}{3}, 1]$. Such preferences are consistent with Ellsberg behaviour in the static three-colour urn problem, and are dynamically inconsistent.

As I noted above, NW's discussion of this example in §3.2 is a bit too informal. Unfortunately, this lack of precision conceals certain aspects of the problem that are, in my view, quite crucial. I will flesh out their discussion, using the language introduced in Section 2. For simplicity, "time 1" refers to the time when the DM chooses between L and R , and "time 2" refers to the time when, if necessary, the DM chooses between u and d .

As noted in §2, every plan of action maps to a Savage act in a natural way. NW adopt the *reduction* principle, also discussed in §2, and we shall do so here as well. This induces a preference ordering over plans at time 1. In particular, we can say that

$$(L, d) \sim (R, d) \succ (L, u) \sim (R, u).$$

But what do these time-1 preferences over plans actually *mean*? Strictly speaking, the DM only chooses between L and R at time 1: she *does not* also choose between u and d . If she did, then Figure 1 would not be a correct depiction of the situation. Again, in the problem under consideration, the choice between u and d is made at time 2, not at time 1, and right now we are concerned solely with the DM's time-1 preferences.

This seemingly minor quibble is actually the key to reconciling dynamically consistent *behaviour* with dynamically inconsistent *preferences*. A plan of action such as (L, u) can be interpreted as a **subtree** of the tree in Figure 1, in which the *only* choice available at time 1 is L , and the *only* choice available at time 2 is u . Of course, the resulting subtree does not depict a very interesting "decision" problem: no actual choice need be made. But that's exactly the point: *a plan of action ensures commitment to a specific set of contingent choices*. By way of contrast, no such explicit commitment ability is available in the tree of Figure 1 – not at time 1, and not at time 2 following either time-1 action.

Thus, the DM's time-1 "Ellsbergian" preference for the plans (L, d) and (R, d) over the plans (L, u) and (R, u) translates to the following statement:

If the DM could commit to a specific time-2 choice, then at time 1 she would prefer to commit to d rather than u . This is true regardless of the DM's planned choice of L vs. R .

The important point here is that the DM's time-1 preferences for *commitment to d* over *commitment to u* are logically quite distinct from her *actual time-2 preferences* over *u* vs. *d*. In fact, dynamic consistency is precisely the *assumption* that time-1 "commitment preferences" coincide with actual time-2 preferences; if these objects were not logically distinct, dynamic consistency would be a tautology.

We are thus led to consider the DM's actual time-2 preferences. Consider first the easy case of preferences at the decision node following the time-1 choice of *L*. Again, NW are a bit loose here, but the spirit of their discussion is right. As noted in Section 2, the actions *u* and *d* at that node map to the Savage acts (10, 0, 10) and (0, 10, 10) respectively, just like the time-1 plans (*L, u*) and (*L, d*). Furthermore, the DM receives no information between time 1 and time 2 if she chooses *L*. Hence, following *L*, the DM is essentially still using her time-1 preferences over acts, and the reduction assumption "buys" us dynamic consistency: the Ellsbergian DM will continue to strictly prefer *d* to *u*.

Another way to interpret this is to say that, if the DM chooses *L* at time 1 in the tree of Figure 1, and wishes to follow *L* with *d*, the additional commitment power afforded by the plan (*L, d*) (that is, the commitment ability she would have if we simply removed action *u* following *L*) is not valuable: at time 2, this DM will be able to carry out her intended action *d* even if, so to speak, her hands are not tied.

Now consider the more interesting case of the DM's preferences at the node following the time-1 choice of *R*. Again apply reduction, and recall that now *u* and *d* now map to the "partial" Savage act (10, 0, *) and (0, 10, *) respectively. For the running example of MEU preferences and prior-by-prior updating, this implies that the DM strictly prefers *u* to *d*.

We are now getting to the heart of the matter. We noted above that this DM strictly prefers (*R, d*) to (*R, u*); however, as we just found out, once she chooses *R*, she then strictly prefers *u* to *d*. Her time-1 preferences over commitment to *u* vs. *d* following *R* are different from her actual time-2 preferences. Again, there is no logical inconsistency here, because we are talking about preferences over different objects. However, Dynamic Consistency is clearly violated.

It is worth reminding the reader that we imagine that the DM has *already* been given a choice to reconsider her preferences. Thus, despite NW's editorializing, there is no further opportunity for her to feel "ashamed" of her preferences.

Let's now relate these preferences to sophistication and the value of information. A sophisticated DM will reason as follows:

I anticipate that, if I choose *R* at time 1, I will then follow *R* with *u*: there is no way I will be able to choose *d* after *R*, even though right now I would very

much like to. On the other hand, I *anticipate* that, if I choose L at time 1, I will follow L with d .

Hence, if I choose L in the tree of Figure 1, it is *as if* I had the opportunity to choose the plan of action (L, d) ; if instead I choose R , it is *as if* I was choosing the plan of action (R, u) .

Since, at time 1, I strictly prefer (L, d) to (R, u) , I should choose L .

Strotz (1955–6, p. 173) said it best: the DM chooses “the best plan among those [s]he will actually follow”. To elaborate, the sophisticated DM correctly anticipates her time-2 preferences, and incorporates them as an “implementation” constraint into her time-1 decision problem.

The action R causes the choice between u and d to be made after learning that the ball drawn is black or red, whereas the action L conveys no further information. Thus, the sophisticated DM under consideration *seemingly* chooses not to receive information by opting for L at time 1. NW follow an illustrious, but still (in my view) unfortunate tradition⁵ in referring to this as a manifestation of “information aversion”.

3.2 The trade-off between value of information and value of commitment

However, the more detailed analysis of the tree in Figure 1 provided above reveals that this conclusion is not warranted. The noted choice of L is *inconclusive* as regards the DM’s attitudes toward information. Moreover, at best, the tree in Figure 1 is too simple to gain a full understanding of information acquisition under ambiguity.

The reason is that the actions L and R differ in *two* dimensions. It is certainly true that, as noted above, R causes the choice between u and d to be made after learning that the ball drawn is black or red, whereas the action L conveys no further information. However, R also causes the choice between u and d to be made on the basis of the DM’s *updated* preferences over acts, whereas the action L ensures that the corresponding choice will be made based upon the DM’s *prior* preferences over acts. Thus, by choosing L , the DM can be certain that, come time 2, she will be able to implement the plan that she deems optimal at time 1; equivalently, action L enables the DM to *commit* to her intended course of action. On the other hand, by choosing R , the DM recognizes that she will not be able to do so – and indeed she will end up choosing an action that, from the perspective of time 1, is undesirable.

NW are rather fond of invoking “basic economic reasoning” throughout their critique, so I shall do the same here. Like any good economic agent, our DM must *trade off* the superior information provided

⁵ See Machina (1989), Epstein and Le Breton (1993), Hanany and Klibanoff (2007b), among others.

by R with the ability to commit to her intended plan afforded by L . A priori, this trade-off could be resolved either way; for the preferences under consideration, it so happens that commitment is the overriding concern, and this leads to the choice of L at time 1.

Of course, if the DM was dynamically consistent, there would be no trade-off, because commitment would simply not be an issue regardless of whether or not information is obtained. But, once again, recall that we are taking dynamic inconsistency as given: when we arrive on the scene of the (alleged) crime, the DM has already been given the opportunity to reconsider her preferences, and has chosen not to do so.

Because the choice of L over R reflects a trade-off between the value of information and the value of commitment, it cannot be taken as evidence of “information aversion”. An analogy with risk attitudes may be useful. Consider the random variables X and Y , where X yields the prizes 10 and 0 with equal probability, and Y yields 2 for sure. Suppose the DM prefers X to Y . Could we then conclude that this choice is driven by risk *appeal*? After all, X is risky and Y is certain! The answer is clearly negative, because the DM may simply be trading off risk and return; for instance, a risk-averse DM with square-root utility prefers X to Y . In order to ascertain the DM’s risk attitudes, we must ask her to rank every random variable X with the degenerate random variable that delivers the prize $E[X]$ for sure – that is, ask the DM to compare random variables that only differ in their riskiness.

Similarly, we cannot conclude that the DM’s choice of L in Figure 1 reveals “information aversion”. In order to ascertain whether, in fact, this DM is averse to information, we would have to consider a different decision problem, in which the options only differ in their informational content. For instance, we could ask the DM to choose one of the four *plans of action* (L, u) , (L, d) , (R, u) and (R, d) at time 1. If the DM expressed a strict preference for, say, (L, d) over (R, d) , then we could conclude that she *intrinsically* dislikes information. But, again, these are *not* the choices offered in the tree of Figure 1.

In fact, it is straightforward to see that the *intrinsic* value of information, once obstacles to commitment are removed, is necessarily non-negative for any DM with complete and transitive preferences. Specifically, if the DM can *commit* to any plan of action, then in particular she can commit to plans wherein the choice of time-2 action is not contingent upon the information received; on the other hand, she can also commit to plans that induce information-contingent actions. Therefore, information expands the set of feasible choices for the DM, and hence it must be valuable, or at worst neutral, if the opportunity to commit is provided.

This argument breaks down in the case of Figure 1 precisely because the initial action R does *not* afford the ability to commit. In this example, because of dynamic inconsistency, information does not expand the set of feasible plans; consequently, there is nothing “embarrassing” in the DM’s choice not to avail herself of information in this particular problem.

In other words, when the analysis is carefully done, ambiguity and dynamic inconsistency do not subvert any deep “economic principle”. On the contrary, basic economic intuition plays a key role in understanding the subtleties in dynamic choice under ambiguity.

3.3 A limitation of NW’s example, and an alternative

Finally, as noted above, the tree in Figure 1 is much too simple to illustrate the issues related to information acquisition under ambiguity. In particular, in the decision problem under consideration, *information is intrinsically worthless*, even assuming that the DM can commit.

This is easy to see: the plans (L, u) and (R, u) map to the same Savage act $(10, 0, 10)$, and similarly the plans (L, d) and (R, d) both map to the act $(0, 10, 10)$. Thus, in this example, information does *not* expand the set of feasible Savage acts. Hence, it is worthless. For this reason, *information would be worthless even if the DM was dynamically consistent!* Thus, NW’s choice of example is somewhat unfortunate.

We can obtain a slightly more interesting example by modifying the tree in Figure 1 as follows. First, replace the payoff 10 in state y following L and d with 0. Second, replace the payoff 10 in state y following R with a choice between u , which yields a payoff of 10, and d , which yields a payoff of 0. Thus, in this modified tree, the plan (L, u) maps to the act $(10, 0, 10)$, and the plan (L, d) maps to the act $(0, 10, 0)$. More importantly, the plans beginning with the choice of R are now (R, u, u) , (R, u, d) , (R, d, u) and (R, d, d) ; for instance, (R, u, d) means “choose R at time 1, then u at time 2 if you learn that the ball is either black or red, and d otherwise”. We do not have corresponding acts beginning with L , because L intuitively corresponds to a situation in which no information is revealed.

As may be expected, information in this modified tree *is* intrinsically valuable. Again, suppose that the DM can choose among all plans of action (i.e. she can commit to any pair of contingent choices at time 2). Of course, the plans (R, u, u) and (R, d, d) map to the same Savage acts as (L, u) and (L, d) , so the DM can do at least as well by choosing R than by choosing L . However, now she can do strictly better: in particular, the plan (R, d, u) maps to the Savage act $(0, 10, 10)$, which *cannot* be obtained by choosing L at time 1. Thus, this DM assigns strictly positive value to information.

Yet, it turns out that the DM’s inability to commit exactly offsets the value of information. In particular, it is still the case that, upon choosing R and observing that the ball drawn was either black or red, the DM will strictly prefer u to d . Of course, the DM will also strictly prefer u to d if she learns that the ball drawn was yellow. Arguing as we did above, the DM will then conclude that choosing R at time 1 is tantamount to “committing” to the plan (R, u, u) , which is just as good as her no-information optimal plan (L, u) .

In Siniscalchi (2009), I present a slightly richer example (involving four states, but still MEU preferences and prior-by-prior updating) wherein the DM's inability to commit diminishes, but does not eliminate or, worse, more than compensate for the intrinsic value of information. Taken together, these examples suggest that the interplay between ambiguity and information acquisition is subtle, interesting, and in full accordance with basic economic intuition.

3.4 Commitment under ambiguity

NW do acknowledge that a desire for commitment may explain the choice of *L* vs. *R*. However, they immediately go on to assert that commitment cannot be valuable in this setting. In light of the analysis just provided, this statement should sound a bit bizarre. However, upon a careful reading of §3.2 in NW, it becomes apparent that their discussion is only seemingly directed at the issue of "commitment": NW are really objecting to the very possibility that dynamic inconsistency may arise in connection with information acquisition.

NW compare commitment in a situation such as the tree in Figure 1 with commitment in a dynamic game, as well as in a single-person decision setting under "temptation". I want to preserve a fully decision-theoretic perspective on ambiguity, so I fully agree with NW when they assert that justifications for strategic precommitment in games simply cannot apply to the single-person setting under consideration. This leaves the comparison with the temptation literature as the key issue. In this respect, NW begin by observing that, in that literature,

the source of temptation is psychological urges that have an independent motivation. For example, addiction to cigarettes or alcohol is, presumably, founded in the physiology of the brain and thus represents an objective and independently motivated constraint.

For this reason, NW argue, the DM should not feel embarrassed if she tries to precommit against such urges. On the other hand,

[t]he desire for commitment under ambiguity lacks such motivations. A subjective ambiguity representation captures the decision maker's model of his environment. While introspection is unlikely to eliminate physiologically induced urges, [...] the subjective decision model is an entirely different matter. It is a mental construct the decision maker created to help him coherently think about the uncertainty he faces, interpret information, and make decisions. The decision maker can change his model if, upon introspection, he finds it wanting or inadequate.

There are two serious flaws with this argument. The first is that NW are equating the "subjective ambiguity representation" with addiction. The parallel they seem to suggest is the following: once the DM decides to

adopt a “model” or “mental construct” (such as the MEU representation with prior-by-prior updating) to guide her preferences and choices, she becomes so “addicted” to it that, upon receiving information, she simply *must* continue to use it, and thereby make a-priori undesirable choices – like a smoker who simply cannot quit, no matter how hard he tries. Anticipating this, the DM may value commitment – like a smoker who intentionally flushes away his cigarettes.

This parallel strikes me as a bit disingenuous. It is *preferences*, not “subjective ambiguity representations”, that we are interested in. If the DM tentatively decides to adopt a given decision model and updating rule, nothing prevents her from “flushing away” the model if she simply *does not like* its prescriptions! Again, as theorists, and at least for the purposes of the present discussion, we are not interested in the DM’s “tentative” preferences; we should grant as much time as the DM needs to reassess her choices, and then devote our attention to preferences that we can assume to be introspection-proof. Thus, there is no parallel between functional representations of preferences and addiction.

The second flaw is the attempt to draw a distinction between the urge to smoke or imbibe, which they deem “objective and independently motivated”, and the disposition to choose u over d upon learning that the ball drawn was black or red. I simply cannot see how one can usefully and *objectively* distinguish between these phenomena. At the most basic, “hardware” level, both are ultimately driven by physiological processes in the brain, so in this sense they are both “objective”. As for “independence”, I am not sure what that means. Perhaps NW wish to claim that, unlike the urge to smoke, the preference for u vs d is “dependent upon” adopting a given decision model; however, as I have noted above, if the DM was displeased with that preference, she could and should simply “flush away” the model that generated it. Or, perhaps NW wish to suggest that, unlike the DM’s preference for u vs. d , the urge to smoke is triggered “independently” of her rational thought processes; but this begs the question of how one decides which thought processes are rational. In fact, I am a little troubled by any attempt to discriminate between “legitimate” urges, which one must accept as a fact of life but can fortunately be controlled via precommitment, and “illegitimate” urges, which are “unreasonable” and must therefore simply be stamped out.

Finally, recent progress in decision theory and menu choice in particular has clarified why we consider actions such as flushing away cigarettes as the expression of rational behaviour. We can view such actions as optimal choices in a suitably rich environment, as in Gul and Pesendorfer (2001, 2005) and the ensuing literature: the DM strictly prefers a “menu” consisting of mints alone to a “menu” consisting of mints and cigarettes because she knows that, if cigarettes are available, then either she will succumb to temptation or she will have to exert costly self-control.

Thus, there is a clear sense in which getting rid of cigarettes is welfare-improving: it is the DM's own preferred course of action.

In Siniscalchi (2009), I show that the desire for commitment in the presence of ambiguity has *exactly the same formal justification*, and the same welfare implications. It is an expression of rationality in an environment wherein the basic objects of choice are suitably rich – in particular, decision trees rather than Savage acts.

In light of these observations, my impression is that NW do not have a compelling, distinct argument against commitment per se: they just think that commitment *should not be an issue* in information-acquisition problems. In other words, they think that the DM should be dynamically consistent. This quote (italics added for emphasis) gives away their prejudice:

Sophisticates who are able to plan for all future contingencies are unlikely to persist in ambiguity aversion *when perceiving their dynamic inconsistency*.

There is nothing wrong with expressing a subjective preference for dynamic consistency, as long as it is recognized as such. It is a manifestation of NW's modelling tastes, not a higher-order rationality principle.

4. RELAXING CONSEQUENTIALISM

As noted in Section 2, a different attempt to reconcile dynamic consistency with ambiguity entails relaxing the property of *Consequentialism* enjoyed by EU preferences. In particular, Hanany and Klibanoff (2007b, 2007a) allow conditional preferences to depend upon the entire decision problem, as well as on the *ex-ante* optimal plan the DM intends to carry out. They characterize the family of “updating rules” that lead to non-consequentialist, but dynamically consistent preferences.

Following NW, I shall focus for definiteness on one distinguished element of the family of dynamically consistent updating rules for MEU preferences, the *ambiguity-maximizing* one (Hanany and Klibanoff, 2007b, Definition 7). In the simple examples considered here and in NW, this rule works roughly as follows. Fix a decision problem, an *ex-ante* optimal plan p , a conditioning event E , and a set of MEU priors C . Consider first those elements of C for which the plan p is *ex-ante* optimal, given the set of feasible plans; call the resulting set $C(p)$. Then, compute the minimum conditional expected utility of p given E over all priors in $C(p)$, and denote it by $u^{\min}(p)$; finally, update all priors in C that assign conditional expected utility not smaller than $u^{\min}(p)$ to p . As is easy to verify, this ensures that the designated *ex-ante* optimal plan will also be optimal conditional upon E . Note that this rule depends crucially upon both the set of feasible plans in the overall decision problem, as well as (in case of ties) upon the optimal plan that the DM wishes to implement.




FIGURE 2. Sunk costs? (NW fig. 2).

While I am perhaps less sanguine than NW as regards departures from Consequentialism, I do share their misgivings. I do not think their “sunk-cost” example is especially compelling, and I will briefly discuss the reason why. I will then propose a related example that highlights a conceptually troublesome implication of non-consequentialist updating: a DM who receives the same information in two different trees may perceive ambiguity differently.

4.1 The sunk-cost example

NW’s example is reproduced in Figure 2. The urn model is as in the preceding section, and we continue to consider the same MEU preferences over acts. Now the DM has the option to Invest S dollars so as to improve the payoff in case a yellow ball is drawn. Subsequently, in case the ball drawn is not yellow, she can choose u or d , which correspond to betting on black or red respectively. The final payoffs are as depicted in Figure 2.

Assume first that $S = 3$. As NW describe in Example 5, the *ex-ante* optimal plan for the DM is (I, d) . However, if the DM were to update prior-by-prior, the optimal choice following I would be u , not d . Instead, the Hanany and Klibanoff ambiguity-maximizing rule prescribes that only priors P with $P(r) \in [\frac{1}{6}, \frac{2}{3}]$ will be updated. As a result, it is conditionally optimal for the DM to follow I with d , as is easily verified. However, if S is high enough, the *ex-ante* optimal plan is (I, u) , and in this case the entire set C will be updated.

NW use this example to make two main points. First, the set of posterior beliefs is “distorted” so as to ensure that the *ex-ante* optimal plan will be implemented. Second, as must be the case whenever *ex-ante* preferences are consistent with the Ellsberg paradox and the agent

is dynamically consistent (cf. NW's example 2 on p. 14), conditional preferences depend upon the "sunk cost" S – if the DM chooses to pay the sunk cost, her conditional decision changes.

However, the fact that conditional beliefs and preferences may be "context-dependent" is the very *definition* of the non-consequentialist approach. "Distorting" beliefs so as to ensure implementation of the *ex-ante* optimal plan is not merely one particular aspect of non-consequentialist dynamic choice under uncertainty: it is its very *raison d'être*, the natural counterpart of the approach taken by Machina (1989) in the context of risky choice. A modeller may not like this approach, but this is only a reflection of his or her own tastes: there is nothing inherently or logically inconsistent in the Hanany and Klibanoff updating rules.

Yet, NW's example points to a more subtle issue.⁶ While there is nothing internally inconsistent in the fact that conditional beliefs could be "context-dependent", it is important to emphasize that *this cannot happen if preferences are probabilistically sophisticated*. In particular, Theorem 3 in Machina and Schmeidler (1992), and results in Epstein and Le Breton (1993, esp. p. 10) imply that, for such preferences, it is possible to define conditional preferences that are non-consequentialist, but agree on the relative likelihood of events regardless of the "context". In other words, even granting that consequentialism may be violated, the phenomenon of belief "distortion" is unique to settings characterized by ambiguity. In my view, this provides a rationale for *maintaining* consequentialism, if one is interested in ambiguity-sensitive behaviour.

Finally, I am of two minds regarding the reference to sunk costs. I can see the rhetorical benefits of eliciting a knee-jerk reaction from economists; however, the point is that the *exact same conclusions about conditional preferences* could be drawn if the problem in Figure 2 is modified by setting S to zero, and considering behaviour in the subtrees beginning with I and $\neg I$ separately. The sunk cost merely determines which plan of action is *ex-ante* optimal, and hence which initial action is taken: it has no *direct* effect on conditional preferences.

4.2 A different example

Since, as I noted above, I am sympathetic to the spirit of NW's comments regarding departures from consequentialism, I provide an alternative example that, in my opinion, makes a somewhat more forceful point. Loosely speaking, relaxing consequentialism so as to guarantee dynamic consistency leads to behaviour that is at odds with our very understanding of ambiguity. The example is taken from Siniscalchi (2009).

⁶ Indeed, I was first made aware of this point by Al Najjar!




FIGURE 3. A dynamic version of the Ellsberg Paradox; $x \in \{0, 1\}$.

The state space is again $\{b, r, y\}$, with the same interpretation as in Figure 1; the composition of the urn will be described below. As in the preceding section, assume that the DM initially has MEU preferences, is risk-neutral for simplicity, and updates her beliefs prior-by-prior upon learning that the ball drawn is not yellow. The decision problem depicted in Figure 3 is parameterized by $x \in \{0, 1\}$: call f_0 (resp. f_1) the tree in which $x = 0$ (resp. $x = 1$). At the initial node (time 1), the DM can Stop and obtain a payoff of 0, or Continue, in which case she learns whether the ball drawn was yellow. If it was, she gets x . Otherwise, she gets to bet on Black or Red.

Suppose that the urn contains 30 black balls, and no more than 30 red balls. The DM’s prior MEU preferences are then characterized by the set C' of probabilities P such that $P(b) = \frac{1}{3}$ and $P(r) \leq \frac{1}{3}$; note that C' contains the uniform prior on $\{b, r, y\}$, denoted P_u . The plan (C, R) is *ex-ante* optimal in the tree f_1 ; to ensure dynamic consistency, the Hanany and Klibanoff ambiguity-maximizing rule requires that only the uniform prior P_u be updated upon reaching the second decision node; all other elements of C' must be discarded. Therefore, the DM must have EU preferences upon learning that the ball drawn was not yellow in the tree f_1 . By way of contrast, the unique optimal plan in the tree f_0 is (C, B) , and the Hanany and Klibanoff ambiguity-maximizing rule prescribes that all priors in C' be updated: thus, the DM has non-degenerate MEU preferences at the second decision node in the tree f_0 .⁷ Hence, conditional upon $\{b, r\}$, the DM *does not* perceive ambiguity in f_1 , but *does* perceive it in f_0 , despite the fact that she receives the same information.

This conclusion stands in sharp contrast with the prevailing view in the literature, which emphasizes that ambiguity is an informational phenomenon. For instance, Ellsberg defines ambiguity as “a quality

⁷ The example can be modified so that (C, R) is the *only ex-ante* optimal plan in f_1 : choose $\epsilon > 0$ and let $C'(\epsilon)$ contain all priors P' such that $P'(b) = \frac{1}{3}$ and $P'(r) \leq \frac{1}{3} + \epsilon$. Now, in the tree f_1 , the Hanany and Klibanoff rules update only priors $P \in C'(\epsilon)$ with $\frac{1}{3} \leq P(r) \leq \frac{1}{3} + \epsilon$, so that, for ϵ small, the DM is “approximately” ambiguity-neutral.

depending on the amount, type, reliability and ‘unanimity’ of the information” (Ellsberg 1961: 657).⁸ If the non-consequentialist approach advocated by Hanany and Klibanoff is adopted, ambiguity must be allowed to depend upon the entire “context” of the (conditional) decision problem, including foregone payoffs and the *ex-ante* optimal plan that is to be carried out. This strikes me as a significant departure from our usual understanding of ambiguity, and a high price to pay in exchange for dynamic consistency.

One may object that the “information” the DM receives at the second decision node in Figure 3 includes the fact that she has avoided the payoff 0 in the tree f_0 , and foregone the payoff 1 in the tree f_1 . But doing so would broaden the interpretation of “information” much beyond its conventional meaning of “partial resolution of the uncertainty” (as evinced from expressions such as “information partition”). After all, there is no uncertainty about the payoff that *might* have been obtained if the ball drawn had been yellow.

I emphasize that the above arguments are specific to choice in the presence of ambiguity. They clearly do not apply to the setting of *risky choice*: if probabilities are “given”, ambiguity simply cannot arise. Second, even in the setting of choice under uncertainty, these arguments do not apply to non-EU preferences that are *probabilistically sophisticated*: in this case, by definition the DM does not perceive any ambiguity.

5. CONCLUSION

To summarize, in my view NW do not provide a compelling argument indicating that dynamic behaviour under ambiguity is riddled with insurmountable difficulties and inconsistencies. Rather, they provide (mostly) well-argued rationales for *their own* modelling choices: simply stated, they are not willing to give up either Consequentialism or Dynamic Consistency, or even relax them in part. However, known results imply that a modeller with these modelling (meta) preferences cannot incorporate ambiguity in his or her analysis.

As a partial remedy, NW suggest that Ellsbergian preferences may be the result of misapplied heuristics. Subjects incorrectly expect the experimenter to be “out to get them”, and behave as if they were playing a zero-sum game. However, I find this suggestion unsatisfactory. As an empirical matter, the amount of experimental evidence confirming Ellsberg’s observations is, by now, rather substantial. It is hard to believe that the vast majority of experimental subjects, across several countries,

⁸ See also Epstein and Zhang (2001: 265), Gilboa and Schmeidler (1989: 142), Schmeidler (1989: 571), Ghirardato, Maccheroni and Marinacci (2004: 134), Klibanoff, Marinacci, and Mukerji (2005: 1849).

age groups, backgrounds, etc., are *all* victims of “misapplied heuristics”, or (worse) paranoia. Moreover, some recent experiments actually attempt to control for possible misapplied heuristics: see for instance Hey, Lotito, and Maffioletti (2008).

Ultimately, however, I think NW’s critique can be interpreted constructively by proponents of ambiguity. NW’s paper does show that it is difficult to debate the appeal of different approaches to dynamic choice under ambiguity from a purely abstract (“normative”) point of view. New empirical and experimental evidence concerning how individuals *actually* behave in dynamic situations under ambiguity may provide more effective guidance for theoretical development in this exciting field.

REFERENCES

- Al-Najjar, N. and J. Weinstein. 2009. The ambiguity aversion literature: a critical assessment. *Economics and Philosophy*, 25.
- Ellsberg, D. 1961. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics* 75: 643–69.
- Epstein, L. and M. Le Breton. 1993. Dynamically consistent beliefs must be Bayesian. *Journal of Economic Theory* 61: 1–22.
- Epstein, L. and M. Schneider. 2001. Recursive multiple-priors. *Journal of Economic Theory* 113: 1–31.
- Epstein, L. and J. Zhang. 2001. Subjective probabilities on subjectively unambiguous events. *Econometrica* 69: 265–306.
- Ghirardato, P. 2002. Revisiting Savage in a conditional world. *Economic Theory* 20: 83–92.
- Ghirardato, P. F. Maccheroni and M. Marinacci. 2003. Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory* 118: 133–73.
- Gilboa, I. and D. Schmeidler. 1989. Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics* 18: 141–53.
- Gul, F. and W. Pesendorfer. 2001. Temptation and self-control. *Econometrica* 69(6): 1403–35.
- Gul, F. and W. Pesendorfer. 2005. The Revealed Preference theory of changing tastes. *Review of Economic Studies* 72(2): 429–48.
- Hammond, P. 1988. Consequentialist foundations for expected utility. *Theory and Decision* 25: 25–78.
- Hammond, P. 1989. Consistent plans, consequentialism, and expected utility. *Econometrica* 57: 1445–9.
- Hanany, E. and P. Klibanoff. 2007a. *Updating uncertainty averse preferences*. mimeo, Kellogg Graduate School.
- Hanany, E. and P. Klibanoff. 2007b. Updating preferences with multiple priors. *Theoretical Economics* 2(3): 261–98.
- Hey, J., G. Lotito and A. Maffioletti. 2008. *The descriptive and predictive adequacy of theories of decision making under uncertainty/ambiguity*. Technical report, Department of Economics, University of York.
- Karni, E. and Z. Safra. 1989. Dynamic consistency, revelations in auctions and the structure of preferences. *Review of Economic Studies* 56: 421–33.
- Karni, E. and D. Schmeidler. 1991. Atemporal dynamic consistency and expected utility theory. *Journal of Economic Theory* 54: 401–8.
- Klibanoff, P., M. Marinacci and S. Mukerji. 2005. A smooth model of decision making under ambiguity. *Econometrica* 73(6): 1849–92.

- Kreps, D. and E. Porteus. 1978. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica* 46: 185–200.
- Machina, M. 1989. Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature* 27: 1622–68.
- Machina, M. and D. Schmeidler. 1992. A more robust definition of subjective probability. *Econometrica* 60: 745–80.
- Savage, L. 1954. *The foundations of statistics*. New York: Wiley.
- Schmeidler, D. 1989. Subjective probability and expected utility without additivity. *Econometrica* 57: 571–87.
- Segal, U. 1990. Two-stage lotteries without the reduction axiom. *Econometrica* 58: 349–77.
- Siniscalchi, M. 2009. *Dynamic choice under ambiguity*. Technical Report 1430, Northwestern University CMS-EMS Discussion Paper.
- Strotz, R. 1955–56. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies* 23: 165–80.