# Dynamic Semi-Consistency

Sophie Bade*

February 10, 2022

**Abstract**

Semi-consistent conditional preferences are inconsistent enough for different ambiguity attitudes to manifest themselves in different behavior and consistent enough for information to be generically valuable. To simultaneously achieve these to desiderata I assume exactly one type of dynamic inconsistency: agents do not update their preferences upon learning independent randomization outcomes.

Keywords: Ambiguity Aversion, Dynamic Consistency, Games with Incomplete Information. *JEL Classification Numbers*: C71, D81, D82.

*Royal Holloway College, Universiy of London and Max Planck Institute for Research on Collective Goods, Bonn sophie.bade@rhul.ac.uk I would like to thank Jürgen Eichberger, Shiri Elon, Gabi Gayer, Michael Mandler, and Larry Samuleson.

# 1 Introduction

Does the behavior of ambiguity averse agents differ from that of expected utility maximizers? The answer depends on the agents' dynamic choice behavior. If agents are dynamically consistent and have access to any randomization device, then all behavior looks alike - ambiguity attitudes do not matter. For different ambiguity attitudes to translate to different behavior, some dynamic inconsistencies must be assumed. Dynamically inconsistent agents, however, may prefer to remain ignorant. Faced with this choice, I suggest a compromise: A dynamically semi-consistent agent is consistent with respect to signals about the world. He does, however, not update his preference over acts if he learns the outcome of a randomization device he considers as independent of these acts.

The Raiffa [32] critique of the Ellsberg paradox illustrates the shortcomings of dynamic consistency. An agent may bet on the color of a ball drawn from an urn that contains 10 black and yellow balls in unknown proportion. The bets $b$ and $\bar{b}$ respectively pay 1 if a black or yellow ball is drawn and nothing otherwise. The Ellsberg paradox consists in the agent's strict preference of a lottery $q$ yielding $1-\epsilon$ with probability one half and nothing otherwise. An expected utility maximizer must prefer $b$ or $\bar{b}$ (or both) to $q$. Raiffa [32] argues that even an ambiguity averse agent who prefers $b$ and $\bar{b}$ to $q$ should not choose $q$. Instead, he should condition his choice of $b$ or $\bar{b}$ on the toss of a fair coin to ensure a winning probability of one half - no matter the proportion of black and yellow balls in the urn. Dynamic consistency is crucial: Raiffa's [32] argument assumes that the agent follows his ex-ante optimal plan and indeed chooses the bet prescribed by the coin-toss. Theorem 3 generalizes this observation to a large set of decision problems: any choice of a dynamically consistent agent (with any ambiguity attitude) is observationally equivalent to the choice of an expected utility maximizer.

For an agent's ambiguity aversion to become visible in the Ellsberg paradox, the agent's behavior must be at least somewhat dynamically inconsistent. If the agent's original preference over $b$, $\bar{b}$, and $q$, for example, coincides with the agent's preference over these bets conditioning on the coin toss, then he will not follow the plan laid out by Raiffa [32]. Unfettered dynamic inconsistency, however, yields counterintuitive models of learning as forgetting as illustrated by the following example.

Say an agent may either bet on an event $\rho$, occurring with probability $\frac{3}{5}$, or its complement $\lambda$. If the agent wins he obtains 1, otherwise he gets nothing. The agent may also choose the above lottery $q$, which yields $1-\epsilon$ with probability one half and nothing otherwise. Ex ante the agent strictly prefers to bet on $\rho$. Now use the above urn filled with black and yellow balls to generate ambiguous signals. The agent receives the signal $L$ if either $\rho$ and yellow or if $\lambda$ and black. Otherwise he receives the signal $R$. Suppose the agent learns $L$ before his choice. Since the agent does not know the number of black balls in the urn, he does not know the conditional probability of $\rho$ given $L$. Conditionally on the signal $L$ the agent's belief on $\rho$ is ambiguous. Being ambiguity averse, the agent may, upon learning $L$, prefer $q$ to either of the two bets. The same reasoning applies to the signal $R$. So if the agent considers the urn ambiguous enough, he chooses $q$ for either signal. In this problem, the signals $L$ and $R$ induce the agent to forget information he already had: the agent goes from the belief that $\rho$ occurs with probability $\frac{3}{5}$, to a vague idea about this probability for either signal. Since always choosing $q$ is ex-ante inferior to betting on $\rho$, the agent is better off ignoring the signals $L$ and $R$. Why would this agent pay any attention to the signals, given that they lead him to make worse choices?

The notion of dynamic semi-consistency addresses the issues illustrated above examples by retaining enough consistency for informative signals to be valuable while permitting enough inconsistency for ambiguity attitudes to matter.

The two examples revolve around two - orthogonal - types of uncertainty: randomization outcomes (coin-toss) and signals about the underlying state ($L$ or $R$). My notion of dynamic semi-consistency distinguishes between these two sources of uncertainty. While I impose dynamic consistency with respect to signals a different normative criterion, namely independence, governs conditional preferences given randomization outcomes. For a Bayesian agent the two principles generate the same conditional preferences: for any outcome of an (independent) randomization device the agent's prior and posterior on payoff relevant events coincide. The Ellsberg example demonstrates that -

without the assumption of Bayesianism - the two principles may clash. Dynamic consistency requires that the agent who sets out the ex-ante optimal plan to bet on black if heads and to bet on yellow if tails updates his preferences so that betting on black and yellow are respectively best if the coin comes up heads or tails - contradicting the independence of the coin.

Now consider the choice of a semi-consistent agent in the Ellsberg paradox, assuming that he ex-ante prefers the lottery $q$ to betting on black or yellow. Since the coin is independent of the color composition of the urn, such an agent does not update his preference over the bets and $q$ upon learning the outcome of the coin-toss. So the agent knows that he won't follow through with the ex-ante optimal plan to bet on black if heads and to bet on yellow otherwise. As the agent prefers $q$ to the bets whether the coin comes up heads or not, he has to resign himself to choose $q$ in either case. The ambiguity aversion of a semi-consistent agent is, in sum, visible in the Ellsberg example. Semi-consistency is, at the same time, restrictive enough to rule out the pathological second example. Being dynamically consistent with respect to signals, a semi-consistent agent will not choose the ex-ante inferior option $q$ for every signal.

I first show that the model of choice sets as sets of lotteries over *all* actions implicitly makes an assumption of dynamic consistency. If one only assumes dynamic semi-consistency agents only choose from the subset of lotteries over *optimal actions* (Theorem 1). I show next that a (weakly) ambiguity averse agent who is fully dynamically consistent *and* considers randomization devices independent in the sense that he does not update his preference upon learning independent randomization events is Bayesian (Theorem 2). The latter result implies that any model of strictly ambiguity averse agents must either drop the assumption of full dynamic consistency or of independent randomization devices. Faced with this dilemma, I extend the Raiffa [32] critique to a much larger set of problems allowing for informational signals to show that any dynamically consistent choice can be explained as a Bayesian choice (Theorem 3). This observational equivalence holds in the comprehensive class of uncertainty averse preferences characterized by Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio [6] which of course contains Gilboa and Schmeidler's [15] maxmin expected utilites as well as

Klibanoff, Mariniacci and Mukerji's [25] smooth model.

Dynamic semi-consistency then assumes a - small and systematic - relaxation of dynamic consistency. Semi-consistent agents do not update their preferences upon learning the outcome of independent randomization devices but are dynamically consistent with respect to all other events. The Ellsberg paradox demonstrates that the behavior of semi-consistent ambiguity averse agent may differ from that of Bayesians: No Bayesian would choose $q$ for both sides of coin. But what about models where ambiguity only enters via the agent's information, such as the second example above? Theorem 4 shows that, even in this case, the dynamically semi-consistent behavior of ambiguity averse and neutral agents may differ. Theorem 4 is couched within the most well-studied model of ambiguity aversion: preferences not only have a maxmin expected utility representation following Gilboa and Schmeidler [15], following Epstein and Schneider [13], the agent's information is modeled such that dynamic consistency and consequentialism do not conflict.

# 2 Decision Problems

There is a finite set of outcomes $X$, and $\Delta X$ is the set of all lotteries over outcomes.[1] There is a state space $\Omega$ endowed with a $\sigma$-algebra $\Sigma$. Ex ante preferences are defined over $\Sigma$-measurable Anscombe-Aumann acts $F : \Omega \to \Delta X$ mapping states to lotteries over outcomes. An expected utility $u : \Delta X \to \mathbb{R}$ represents the agent's preference over constant acts which map each state $\omega \in \Omega$ to the same lottery in $\Delta X$. The agent's preference over all $\Sigma$-measurable acts $F : \Omega \to \Delta X$ is represented by this utility $u$ and a complete and transitive ranking $\succsim$ on utility valued acts, so that one act $F$

[1] For any finite set $S$, $\Delta S$ denotes the set of all lotteries on $S$.

is weakly preferred to another act $F'$ if $u \circ F \succsim u \circ F'$.[2]

To save on notation, decision problems are defined in terms of utility valued $\Sigma$-measurable acts $f : \Omega \to \mathbb{R}$ with $f : = u \circ F$ for some $\Sigma$-measurable $F : \Omega \to \Delta X$. The preference $\succsim$ is monotonic in the sense that an act $f$ is strictly preferred to a different act $f'$ if $f$ yields a higher utility than $f'$ in every state, formally $f(\omega) > f'(\omega)$ for all $\omega \in \Omega$ implies $f \succ f'$. If $f$ is constant on some event $E$, then $f(E)$ denotes $f(\omega)$ with $\omega \in E$. For any two $\Sigma$-measurable acts $f$ and $f'$ and any $\alpha \in (0, 1)$ define the convex combination $\alpha f + (1-\alpha)f'$ so that $(\alpha f + (1-\alpha)f')(\omega) = \alpha f(\omega) + (1-\alpha)f'(\omega)$ for all $\omega \in \Omega$. There exist $x, x' \in X$ such that $u(x) \neq u(x')$.

Fix a $\sigma$-algebra $\Sigma' \subset \Sigma$ and say $\succsim'$ is the restriction of $\succsim$ to the set of $\Sigma'$-measurable acts. The agent is **Bayesian with respect to** $\Sigma'$ if $\succsim'$ has an expected utility representation $U(f) = \int_{\omega \in \Omega} f(\omega) d\pi'(\omega)$ with some $\Sigma'$-measurable prior $\pi'$. Anscombe and Aumann [1] show that such an expected utility representation exists if and only if $\succsim'$ is continuous and satisfies independence axiom. The preference $\succsim'$ satisfies **independence axiom** if $f \sim' f'$ implies $\alpha f + (1-\alpha)f' \sim' f$ for any two $\Sigma'$-measurable acts $f$ and $f'$ and any $\alpha \in (0, 1)$. The preference $\succsim'$ is **continuous** if for all $\Sigma'$-measurable acts $f$, $f'$, and $f''$ with $f \succ' f' \succ' f''$ there exist numbers $\alpha, \alpha' \in (0, 1)$ such that $\alpha f + (1-\alpha)f'' \succ' f' \succ' \alpha' f + (1-\alpha')f''$. If the independence axiom is not assumed to hold for $\succsim'$, then $\Sigma'$ is **ambiguous**. If $f \sim' f'$ implies $\alpha f + (1-\alpha)f' \succsim' f$ for all $\Sigma'$-measurable acts $f, f'$ and all $\alpha \in (0, 1)$ then $\succsim'$ is **(weakly) ambiguity averse**. A (weakly) ambiguity averse preference $\succsim'$ is **strictly ambiguity averse** if there exist two $\Sigma'$-measurable acts $f, f'$ and an $\alpha \in (0, 1)$ such that $f \sim' f'$ as well as $\alpha f + (1-\alpha)f' \succ' f$.[3]

[2]A preference $R$ on acts $F : \Omega \to \Delta X$ can represented via such a $\succsim$ on utility valued acts if $R$ is complete, transitive, monotonic, risk independent, and risk continuous. The preference $R$ is monotonic if $F(\omega)RF'(\omega)$ for all $\omega \in \Omega$ implies $FRF'$ for any $F, F' : \Omega \to \Delta X$. It is risk independent if $pRq$ implies $\alpha p + (1-\alpha)rR\alpha q + (1-\alpha)r$ for all $\alpha \in (0, 1)$ and all constant acts $p, q, r \in \Delta X$. It is risk continuous if for any sequences $(p_n)_{n=1}^{\infty}$ and $(q_n)_{n=1}^{\infty}$ of constant acts in $\Delta X$ with $\lim p_n = p$ and $\lim q_n = q$ and $p_n R q_n$ for all $n$ we have $pRq$. The restriction of the transitive, complete, risk independent and risk continuous $R$ to constant acts $\Delta X$ has by the von-Neumann-Morgenstern Theorem an expected utility representation $u : \Delta X \to \mathbb{R}$. Since $R$ is monotonic there exists a preference $\succsim$ on the set of all utility valued acts such that $FRF'$ holds if and only if $u \circ F \succsim u \circ F'$.

[3]Weakly but not strictly ambiguity averse preferences satisfy the independence axiom.

To distinguish randomization events I assume that each event in $\Sigma$ can be represented as the intersection of a randomization event and a "normal" event. So we have $\Sigma \colon = \Sigma^r \cap \mathcal{P}$ where the $\sigma$-algebra $\Sigma^r$ models a **universal randomization device** and the finite algebra $\mathcal{P}$ models all other events. The intersection of the two algebras $\Sigma^r \cap \mathcal{P}$ is defined as the set of all events $E \cap E'$ for $E \in \Sigma^r$ and $E' \in \mathcal{P}$.[4] The agent is Bayesian with respect to $\Sigma^r$ with prior $\pi^r$. The agent is indifferent between using the universal device $\Sigma^r$ and objective randomization. To capture this indifference say that $\Sigma^r$ is **weakly independent** of $\mathcal{P}$, in the sense that $f \sim \sum_{i=1}^m \pi^r(E_i) f_i$ holds for any finite set of $\mathcal{P}$-measurable acts $\{f_1, \cdots, f_m\}$ and partition $\{E_1, \cdots, E_m\} \subset \Sigma^r$ of $\Omega$ and act $f$ with $f(\omega) = f_i(\omega)$ if $\omega \in E_i$ for all $\omega \in \Omega$.[5] If the agent is Bayesian then $\Sigma^r$ is weakly independent of $\mathcal{P}$ if and only if the two algebras are stochastically independent.

The agent gets to choose an action $a$ from a finite set of actions $A$. Different choices $a \in A$ induce different (utility valued) acts $g(a, \cdot) : \Omega \to \mathbb{R}$ that are measurable with respect to an **algebra of payoff relevant events** $\mathcal{R} \subset \mathcal{P}$. The agent may condition his choice of an action $a$ on outcomes of the universal randomization device and on his signals, described by an **information algebra** $\mathcal{Q} \subset \mathcal{P}$. The partitions $\mathcal{Q}^P$ and $\mathcal{R}^P$ generate the algebras $\mathcal{Q}$ and $\mathcal{R}$. A priori, the agent does not rule out any signal $\theta \in \mathcal{Q}^P$, in the sense that he strictly ranks any two $\mathcal{Q}$-measurable acts $f, f' : \Omega \to \mathbb{R}$ with $f(\theta^*) \neq f'(\theta^*)$ for some $\theta^* \in \mathcal{Q}^P$ and $f(\theta) = f'(\theta)$ otherwise.[6] If the agent is Bayesian with respect to $\mathcal{Q}$, then the preceding condition holds if and only if the agent assigns positive probability to each event $\theta$ in the information partition $\mathcal{Q}^P$.

A choice problem is summarized by $(\succsim, \mathcal{Q}, g)$ the agent's preference $\succsim$ over $\Sigma$-measurable (utility-valued) acts, the information partition $\mathcal{Q}$, and the function $g$ which maps the agents actions and all states to utilities. If the agent is Bayesian with respect to $\Sigma$ then the decision problem $(\pi, \mathcal{Q}, g)$ is **Bayesian** where $\pi$ is the $\Sigma$-measurable prior on $\Omega$ that represents $\succsim$. The

[4] Alternatively one could define the state space as a space of two component vectors: $\Omega \colon = \Omega^r \times P$ with $E \subset \Omega^r$ and $E' \subset P$ denoting randomization and other events.

[5] Weak independence also entails (a form of) continuity: the agent is indifferent between any two acts $f$ and $f'$ that only differ on some event $E \in \Sigma^r$ with $\pi^r(E) = 0$.

[6] As $f$ and $f'$ are utility-valued, they yield different utilities in the event $\theta$ if $f(\theta) \neq f'(\theta)$.

agent's preference has a maxmin expected utility (**MMEU**) representation following Gilboa and Schmeidler [15] if there exists a convex and compact set $C$ of $\Sigma$-measurable priors $\pi$ on $\Omega$, such that $\succsim$ is represented by $U(f) = \min_{\pi \in C} \sum_{\omega \in \Omega} f(\omega)\pi(\omega)$. Since $\succsim$ is a preference over utility valued acts any MMEU preference is defined by a set of beliefs $C$. An agent with a MMEU preference is Bayesian if $C$ is a singleton, otherwise the agent is strictly ambiguity averse.[7]

In the problem $(\succsim, \mathcal{Q}, g)$ the agent chooses a $\Sigma^r \cap \mathcal{Q}$-measurable **plan** $\mathfrak{p} : \Omega \to A$, which specifies an action for each event the agent may observe. If the plan $\mathfrak{a}$ is measurable with respect to the agent's information algebra $\mathcal{Q}$, then it is as **pure plan** and $\mathcal{A}$ is the set of all such pure plans. For any plan $\mathfrak{p} : \Omega \to A$, define $\mathcal{A}^{\mathfrak{p}}$ as the set pure plans sometimes chosen by $\mathfrak{p}$, so formally $\mathcal{A}^{\mathfrak{p}}$ consists of all pure plans $\mathfrak{a}$ such that $\mathfrak{p}(\omega) = \mathfrak{a}(\omega)$ for all $\omega \in E$ for some non-empty $E \in \Sigma^r$. Any plan $\mathfrak{p} : \Omega \to A$ can then be represented as a partition $\{E^{\mathfrak{p}}(\mathfrak{a})\}_{\mathfrak{a} \in \mathcal{A}^{\mathfrak{p}}} \subset \Sigma^r$ of $\Omega$ where for all $\mathfrak{a} \in \mathcal{A}^{\mathfrak{p}}$ we have $\mathfrak{p}(\cdot) = \mathfrak{a}(\cdot)$ on $E^{\mathfrak{p}}(\mathfrak{a})$. We can interpret $\{E^{\mathfrak{p}}(\mathfrak{a})\}_{\mathfrak{a} \in \mathcal{A}^{\mathfrak{p}}}$ as a particular randomization device that generates the plan $\mathfrak{p}$ by choosing the pure plan $\mathfrak{a}$ in the randomization event $E^{\mathfrak{p}}(\mathfrak{a})$.

Any such plan $\mathfrak{p}$ induces a $\Sigma$-measurable act $g(\mathfrak{p}(\cdot), \cdot) : \Omega \to \mathbb{R}$. Due to the assumption that the universal randomization device $\Sigma^r$ is weakly independent of $\mathcal{P}$, the agent is indifferent between a plan $\mathfrak{p}$ and a lottery over pure plans $p \in \Delta\mathcal{A}$ if both choose any given pure plan with the same probability. Expressed formally, we have $g(\mathfrak{p}(\cdot), \cdot) \sim \sum_{\mathfrak{a} \in \mathcal{A}} p(\mathfrak{a}) g(\mathfrak{a}(\cdot), \cdot)$ if $p(\mathfrak{a}) = \pi^r(E^{\mathfrak{p}}(\mathfrak{a}))$ holds for all $\mathfrak{a} \in \mathcal{A}$. The randomization device $\Sigma^r$ is **rich** in the sense that for any lottery $p \in \Delta\mathcal{A}$ with $\mathcal{A}^* \colon = \{\mathfrak{a} \colon p(\mathfrak{a}) > 0)\}$ there exists a partition $\{E(\mathfrak{a})\}_{\mathfrak{a} \in \mathcal{A}^*} \subset \Sigma^r$ of $\Omega$ with $\pi^r(E(\mathfrak{a})) = p(\mathfrak{a})$ for all $\mathfrak{a} \in \mathcal{A}^*$.

To skirt the issue of conditional preferences given null events, I assume that any pure plan $\mathfrak{a}$ is either adopted with positive probability or never, so

[7] For a MMEU with set of beliefs $C$, the above assumption on the possibility of every signal translates to $0 < \pi(\theta)$ for all $\pi \in C$ and all $\theta \in \mathcal{Q}^P$.

the agent's choice set is[8]

$$P(\mathcal{A})\colon = \{\mathfrak{p} : \pi^r(E^{\mathfrak{p}}(\mathfrak{a})) > 0 \text{ for all } \mathfrak{a} \in \mathcal{A}^{\mathfrak{p}}\}.$$

For any set of pure plans $\mathcal{A}' \subset \mathcal{A}$, the set $P(\mathcal{A}')$ is the set of all plans $\mathfrak{p} \in P(\mathcal{A})$ with $\mathcal{A}^{\mathfrak{p}} \subset \mathcal{A}'$.

# 3 Two Examples

To formalize the two examples in the introduction say $X\colon = \{0, 1-\epsilon, 1\}$ is the set of outcomes and $u(x) = x$ for all $x$, so that the utility of the lottery where the agent wins $1-\epsilon$ with probability $\frac{1}{2}$ is $u(q) = \frac{1-\epsilon}{2}$. The events heads $H$ and tails $T$ are randomization events, so $H, T \in \Sigma^r$. Since the coin is fair we have $\pi^r(H) = \pi^r(T) = \frac{1}{2}$. Conversely, $\mathcal{P}$ contains the events $\rho$ and $\lambda$ as well as the events that a black or yellow ball is drawn, denoted by $\beta$ and $\overline{\beta}$. The bets on black, yellow, $\rho$ and $\lambda$ are respectively denoted $b$, $\overline{b}$, $r$, and $l$. If the agent chooses the action $b$ then he gets $g(b, \omega) = 1$ if $\omega \in \beta$ and $g(b, \omega) = 0$ otherwise. The acts induced by the three other bets $g(a, \cdot)$ for $a \in \{\overline{b}, r, l\}$ similarly yield 1 in the respective winning events $\overline{\beta}$, $\rho$, and $\lambda$ and nothing otherwise. The choice $q$, in contrast, yields a constant expected utility $g(q, \omega) = \frac{1-\epsilon}{2}$ for all $\omega \in \Omega$. In the first example the action set is $\{b, \overline{b}, q\}$ and the partition of payoff relevant events is $\{\beta, \overline{\beta}\}$. In the second example $\{\rho, \lambda\}$ is the partition of payoff relevant events while $\{L, R\}$ with $L\colon = (\rho \cap \beta) \cup (\lambda \cap \overline{\beta})$ and $R\colon = \Omega \setminus L$ is the agent's information partition. The action set in the second example is $\{l, r, q\}$.

In the example of the Ellsberg paradox, the agent does not receive any signals, so his set of actions $\{b, \overline{b}, q\}$ is also his set of pure plans. The randomization device is rich enough to generate any lottery over $\{b, \overline{b}, q\}$. The agent can for example throw a coin and plan to bet on black if heads and to bet on yellow if tails. This plan can be represented as $\mathfrak{p} : \Omega \to A$ with $\mathfrak{p}(\omega) = b$ if $\omega \in H$ and $\mathfrak{p}(\omega) = \overline{b}$ if $\omega \notin H$. Since $\Sigma^r$ is weakly independent

[8]Given that $\Sigma^r$ is rich and weakly independent, this restriction is without loss of generality in the sense that for any plan $\mathfrak{p} \notin P(\mathcal{A})$ there exists a $\mathfrak{p}' \in P(\mathcal{A})$ such that $g(\mathfrak{p}(\cdot), \cdot) \sim g(\mathfrak{p}'(\cdot), \cdot)$ and $\pi^r(E^{\mathfrak{p}}(\mathfrak{a})) = \pi^r(E^{\mathfrak{p}'})(\mathfrak{a})$ for all $\mathfrak{a} \in \mathcal{A}$.

of $\mathcal{P}$ and since $\pi^r(H) = \frac{1}{2}$, the agent is indifferent between $g(\mathfrak{p}(\cdot), \cdot)$ and the objective randomization $\frac{1}{2}g(b, \cdot) + \frac{1}{2}g(\bar{b}, \cdot)$ over the two bets.

In the second example the agent has the non-trivial information partition $\mathcal{Q}^P = \{L, R\}$. We therefore have to consider more complex pure plans, such as the plan $\mathfrak{a} : \Omega \to \{r, l, q\}$ with $\mathfrak{a}(\omega) = r$ if $\omega \in R$ and $\mathfrak{a}(\omega) = l$ otherwise. According to this plan $\mathfrak{a}$ the agent bets on $\rho$ if he learns the signal $R$ and bets on $\lambda$ if he learns $L$. Since there are two signals and three actions ($r$, $l$, and $q$), there are nine different pure plans. The randomization device $\Sigma^r$ contains partitions to generate *any* lottery over these nine pure plans. The agent can, for example, set out a plan $\mathfrak{p}$ where he chooses each pure plan $\mathfrak{a}$ in some event $E^{\mathfrak{p}}(\mathfrak{a})$ that occurs exactly with probability one ninth.

The next section defines dynamically semi-consistent conditional preferences. Semi-consistency requires that preferences over $\mathcal{P}$ measurable acts do not change upon learning independent randomization outcomes. In terms of the above example this assumption yields that the agent in the Ellsberg paradox who prefers $g(q, \cdot)$ to $g(b, \cdot)$ continues to hold this preference upon learning any randomization outcome, in particular $H$ or $T$. Semi-consistency also requires that the agent is dynamically consistent with respect to informative signals. Given that $L$ and $R$ in the second example are such informative signals, this assumption entails that the agent who adopts an ex ante optimal plan $\mathfrak{a}$ specifying actions $\mathfrak{a}(L)$ and $\mathfrak{a}(R)$ for the signals $L$ and $R$ will follow through with these actions upon learning his signal. Theorem 1 then shows that the choice set of a semi-consistent agent can be modelled as the set of lotteries over all his ex ante optimal pure plans. Given that choosing $q$ the unique ex ante optimal pure plan for the agent in the Ellsberg paradox, Theorem 1 yields that a semi-consistent agent must set out of choose $q$ in the Ellsberg paradox.

# 4 Conditional Preferences and Optimal Plans

Fix a choice problem $(\succsim, \mathcal{Q}, g)$ and say that $\mathfrak{p} \in P(\mathcal{A})$ is the plan adopted by the agent. Upon learning events in $\Sigma$ the agent forms conditional preferences. For any preference $\succsim$, plan $\mathfrak{p}$, and event $E \cap \theta$ with $E \in \Sigma^r$, $\pi^r(E) \neq \emptyset$ and $\theta \in \mathcal{Q}^P$ define a **conditional preference** $\succsim^{\mathfrak{p}}_{E\cap\theta}$. This definition has two non-standard aspects. The present conditional preferences are not **consequentialist** as they may depend on the agent's plan $\mathfrak{p}$ (as well as on the fixed decision problem). Conditional preferences are moreover not specified for all (non-null) events but only for non-null randomization events, signals and their intersections. These relaxations allow for dynamic consistency. Machina [30], McClennen [31], Epstein and Le Breton [10] and Ghirardato [14] showed that dynamic consistency implies ambiguity neutrality if consequentialist conditional preferences are defined for all (non-null) events. Conversely, dynamic consistency is compatible with a variety of ambiguity attitudes if conditional preferences are only defined for a select set of events and/or depend on ex-ante plans.[9]

The agent is a **consistent planner** in the sense of Siniscalchi [35]: he only adopts plans that he sticks with. For the agent to stick with a plan $\overline{\mathfrak{p}}$ this plan must for each signal and randomization event prescribe an action that is optimal upon learning both. The set of all such plans is

$$P^*: = \{\mathfrak{p} \in P(\mathcal{A}) : g(\mathfrak{a}(\theta), \cdot) \succsim^{\mathfrak{p}}_{E^{\mathfrak{p}}(\mathfrak{a})\cap\theta} g(a, \cdot) \text{ for all } a \in A, \mathfrak{a} \in \mathcal{A}^{\mathfrak{p}}, \theta \in \mathcal{Q}^P\}.$$

Assuming that the agent does not simultaneously learn signals and randomization outcomes, the chosen plan $\overline{\mathfrak{p}}$ must also satisfy a condition of interim optimality. Upon learning only his signal or the randomization event the plan $\overline{\mathfrak{p}}$ must be best among all plans that he will ultimately follow through with. If the agent learns his signal first, the plan $\overline{\mathfrak{p}} \in P^*$ is interim optimal if $g(\overline{\mathfrak{p}}(\cdot), \cdot) \succsim^{\overline{\mathfrak{p}}}_{\theta} g(\mathfrak{p}(\cdot), \cdot)$ for all $\mathfrak{p} \in P^*$ and all $\theta \in \mathcal{Q}^P$. If the agent learns in the inverse order then $\overline{\mathfrak{p}} \in P^*$ must prescribe optimal pure plans once all randomization uncertainty has resolved. So $g(\mathfrak{a}(\cdot), \cdot) \succsim^{\overline{\mathfrak{p}}}_{E^{\mathfrak{p}}(\mathfrak{a})} g(\mathfrak{a}'(\cdot), \cdot)$ must hold for all chosen pure plans $\mathfrak{a} \in \mathcal{A}^{\overline{\mathfrak{p}}}$ and alternative pure plans $\mathfrak{a}'$.

[9] All upcoming results hold if one permits only one of these two paths out of the dilemma of dynamic consistency. We can either follow the beautiful arguments Hill [20] and assume that there exist consequentialist dynamically consistent preferences for select set of events or we can assume that dynamically consistent preferences depend on ex-ante plans. For the case of MMEU representations Epstein and Schneider [13] developed the theory of consequentialist consistent updating for specific sets of events. Conversely, Hanany and Klibanoff [17] define consistent conditional preferences by letting conditional preferences $\succsim^{\mathfrak{p}}_E$ not only depend on conditioning events $E$ but also on ex-ante choices $\mathfrak{p}$. For an axiomatic approach to updating ambiguity averse preferences that does not impose dynamic consistency see Gilboa and Schmeidler [16].

A $\succsim$-best plan among all plans that the agent will follow through upon learning signals and randomization outcomes (in any order) is an **optimal plan**.

The agent is **(dynamically) consistent** if he follows through with any ex-ante optimal plan. Formally the agent is consistent if

$$\begin{aligned}
&g(\overline{\mathfrak{p}}(\cdot),\cdot) \succsim g(\mathfrak{p}(\cdot),\cdot) \text{ for all } \mathfrak{p} \in P(\mathcal{A}) \Leftrightarrow \\
&g(\overline{\mathfrak{p}}(\cdot),\cdot) \succsim_{\theta}^{\overline{\mathfrak{p}}} g(\mathfrak{p}(\cdot),\cdot) \text{ for all } \mathfrak{p} \in P(\mathcal{A}) \text{ and all } \theta \in \mathcal{Q}^P \Leftrightarrow \\
&g(\mathfrak{a}(\theta),\cdot) \succsim_{E^{\mathfrak{p}}(\mathfrak{a})\cap\theta}^{\overline{\mathfrak{p}}} g(a,\cdot) \text{ for all } a \in A, \mathfrak{a} \in \mathcal{A}^{\overline{\mathfrak{p}}}, \theta \in \mathcal{Q}^P \Leftrightarrow \\
&g(\mathfrak{a}(\cdot),\cdot) \succsim_{E^{\mathfrak{p}}(\mathfrak{a})}^{\overline{\mathfrak{p}}} g(\mathfrak{a}'(\cdot),\cdot) \text{ for all } \mathfrak{a} \in \mathcal{A}^{\overline{\mathfrak{p}}} \text{ and } \mathfrak{a}' \in \mathcal{A}^{\mathfrak{p}}.
\end{aligned}$$

A **semi-consistent** agent is dynamically consistent with respect to signals and considers the randomization device to be strongly independent. The randomization device $\Sigma^r$ is **strongly independent** of $\mathcal{P}$ if the agent does not update his preference over $\mathcal{P}$-measurable acts upon learning any randomization event - no matter which plan he adopted ex ante. Formally, the agent considers the randomization device strongly independent if for any $\mathcal{P}$-measurable acts $f$ and $f'$, ex-ante plan $\overline{\mathfrak{p}} \in P(\mathcal{A})$, $\theta \in \mathcal{Q}^P$, and $E \in \Sigma^r$ with $\pi^r(E) > 0$

$$f \succsim f' \Leftrightarrow f \succsim_E^{\overline{\mathfrak{p}}} f' \text{ and } f \succsim_{\theta}^{\overline{\mathfrak{p}}} f' \Leftrightarrow f \succsim_{E\cap\theta}^{\overline{\mathfrak{p}}} f'.$$

The agent is consistent with respect to signals if for any ex-ante plan $\overline{\mathfrak{p}} \in P^*$, $\mathfrak{a} \in \mathcal{A}^{\overline{\mathfrak{p}}}$, and $E \in \Sigma^r$ with $\pi^r(E) > 0$

$$\begin{gathered}
g(\overline{\mathfrak{p}}(\cdot),\cdot) \succsim g(\mathfrak{p}(\cdot),\cdot) \text{ for all } \mathfrak{p} \in P^* \Leftrightarrow \\
g(\overline{\mathfrak{p}}(\cdot),\cdot) \succsim_{\theta}^{\overline{\mathfrak{p}}} g(\mathfrak{p}(\cdot),\cdot) \text{ for all } \mathfrak{p} \in P^*, \theta \in \mathcal{Q}^P \\
\text{and} \\
g(\mathfrak{a}(\cdot),\cdot) \succsim g(\mathfrak{a}'(\cdot),\cdot) \text{ for all } \mathfrak{a}' \in \mathcal{A} \Leftrightarrow \\
g(\mathfrak{a}(\theta),\cdot) \succsim_{\theta}^{\overline{\mathfrak{p}}} g(a,\cdot) \text{ for all } a \in A, \theta \in \mathcal{Q}^P.
\end{gathered}$$

Theorem 1 shows that it is without loss of generality to represent the agent's choice set as a set of objective lotteries over a set of pure plans. Whether the agent is consistent or semi-consistent there exists a set of pure plans $\overline{\mathcal{A}}$ such that the agent is indifferent between his optimal plan in $P(\mathcal{A})$

and the $\succsim$-best lottery in $\Delta\overline{\mathcal{A}}$. For consistent agents $\Delta\overline{\mathcal{A}}$ equals the set of all possible objective lotteries $\Delta\mathcal{A}$; for a semi-consistent $\overline{\mathcal{A}}$ equals the set of all optimal pure plans $\mathcal{A}^*\colon = \{\mathfrak{a}^* \in \mathcal{A}\colon g(\mathfrak{a}^*(\cdot)) \succsim g(\mathfrak{a}(\cdot)) \text{ for all } \mathfrak{a} \in \mathcal{A}\}$.

**Theorem 1** *Fix a decision problem $(\succsim, \mathcal{Q}, g)$.*

*a) A plan is optimal for a consistent agent if and only if it is indifferent to the $\succsim$-best lottery in $\Delta\mathcal{A}$.*

*b) A plan is optimal for a semi-consistent agent if and only if it is indifferent to the $\succsim$-best lottery in $\Delta\mathcal{A}^*$.*

**Proof** Say $\mathfrak{p}^c$ and $\mathfrak{p}^s$ are optimal plans for a consistent or respectively semi-consistent agent. I first show that $\mathfrak{p}^c$ and $\mathfrak{p}^s$ are $\succsim$-best plans in $P(\mathcal{A})$ and $P(\mathcal{A}^*)$. I then show that any $\succsim$-best plan in $P(\mathcal{A})$ and $P(\mathcal{A}^*)$ is indifferent to a $\succsim$-best lottery in $\Delta\mathcal{A}$ and $\Delta\mathcal{A}^*$.

Since consistent agents follow through with any ex-ante optimal plan, any $\succsim$-best plan $\mathfrak{p}^c$ in $P(\mathcal{A})$ is an optimal plan for a consistent agent.

To see that any $\succsim$-best plan $\mathfrak{p}^s$ in $P(\mathcal{A}^*)$ is an optimal plan for a semi-consistent agent, use first the strong independence of $\Sigma^r$ and then the agent's consistency with respect to signals to see that, for a semi-consistent agent, $P^*$ equals $P(\mathcal{A}^*)$:

$$\begin{aligned}P^* = \{\mathfrak{p} \in P(\mathcal{A}) : g(\mathfrak{a}(\theta), \cdot) \succsim^{\mathfrak{p}}_{E^{\mathfrak{p}}(\mathfrak{a})\cap\theta} g(a, \cdot) \text{ for all } a \in A, \mathfrak{a} \in \mathcal{A}^{\mathfrak{p}}, \theta \in \mathcal{Q}^P\} =\\ \{\mathfrak{p} \in P(\mathcal{A}) : g(\mathfrak{a}(\theta), \cdot) \succsim^{\mathfrak{p}}_{\theta} g(a, \cdot) \text{ for all } a \in A, \mathfrak{a} \in \mathcal{A}^{\mathfrak{p}}, \theta \in \mathcal{Q}^P\} =\\ \{\mathfrak{p} \in P(\mathcal{A}) : g(\mathfrak{a}(\cdot), \cdot) \succsim g(\mathfrak{a}'(\cdot), \cdot) \text{ for all } \mathfrak{a}' \in \mathcal{A}, \mathfrak{a} \in \mathcal{A}^{\mathfrak{p}}\} = P(\mathcal{A}^*).\end{aligned}$$

Since the agent is a consistent planner, $P^*$ must contain all optimal plans of the agent. To see that any $\succsim$-best plan $\mathfrak{p}^s$ in $P(\mathcal{A}^*) = P^*$ is an optimal plan for a semi-consistent agent, we have to show that any such plan $\mathfrak{p}^s$ respects the two conditions of interim optimality. Considering the case that the agent first learns randomization outcomes, fix an outcome $E$ of the randomization device used to generate the plan $\mathfrak{p}^s$, so that there exists a plan $\mathfrak{a} \in \mathcal{A}^{\mathfrak{p}^s}$ with $E = E^{\mathfrak{p}^s}(\mathfrak{a})$, $E \in \Sigma^r$, and $\pi^r(E) > 0$. Since $\mathfrak{p}^s \in P(\mathcal{A}^*)$ and $\mathfrak{a} \in \mathcal{A}^{\mathfrak{p}^s}$, $g(\mathfrak{a}(\cdot), \cdot) \succsim g(\mathfrak{a}'(\cdot), \cdot)$ holds for all $\mathfrak{a}' \in \mathcal{A}$. Since $\Sigma^r$ is strongly independent and since $\pi^r(E) > 0$, $g(\mathfrak{a}(\cdot), \cdot) \succsim^{\mathfrak{p}^s}_E g(\mathfrak{a}'(\cdot), \cdot)$ holds for all $\mathfrak{a}' \in \mathcal{A}$. The plan $\mathfrak{p}^s$ consequently satisfies the condition of interim optimality if the agent first learns the randomization outcome. The agent's dynamic consistency with respect to signals implies that $g(\mathfrak{p}^s(\cdot),\cdot) \succsim_\theta^{\mathfrak{p}^s} g(\mathfrak{p}(\cdot),\cdot)$ holds for all $\mathfrak{p} \in P^*, \theta \in \mathcal{Q}^P$, so that $\mathfrak{p}^s$ also satisfies the condition for interim optimality if the agent first learns his signal.

Since a plan $\mathfrak{p}^c$ ($\mathfrak{p}^s$) is optimal for a consistent (semi-consistent) agent if and only if it is a $\succsim$-best plan in $P(\mathcal{A})$ ($P(\mathcal{A}^*)$), it now suffices to show that the agent is, for any non-empty set of pure plans $\overline{\mathcal{A}}$, indifferent between the best lotteries and plans in $\Delta\overline{\mathcal{A}}$ and $P(\overline{\mathcal{A}})$. So fix any $\emptyset \neq \overline{\mathcal{A}} \subset \mathcal{A}$, $\succsim$-best lottery $\overline{p} \in \Delta\overline{\mathcal{A}}$ and $\succsim$-best plan $\overline{\mathfrak{p}} \in P(\overline{\mathcal{A}})$. Since $\Sigma^r$ is rich there exists a partition $\left(E^{\mathfrak{p}'}(\mathfrak{a})\right)_{\mathfrak{a}\in\mathcal{A}'}$ of $\Sigma^r$ such that $\pi^r(E^{\mathfrak{p}'}(\mathfrak{a})) = \overline{p}(\mathfrak{a})$ for all $\mathfrak{a} \in \mathcal{A}' :\ = \{\mathfrak{a} \mid \overline{p}(\mathfrak{a}) > 0\}$. Since $\Sigma^r$ is weakly independent the agent is indifferent between the plan $\mathfrak{p}' \in P(\overline{\mathcal{A}})$ generated by the randomization device $\left(E^{\mathfrak{p}'}(\mathfrak{a})\right)_{\mathfrak{a}\in\mathcal{A}'}$ and the lottery and $\overline{p}$. The weak independence of $\Sigma^r$ also implies that the agent is indifferent between the plan $\overline{\mathfrak{p}}$ and the lottery $p' \in \Delta(\overline{\mathcal{A}})$ with $p'(\mathfrak{a}) = \pi^r(E^{\overline{\mathfrak{p}}}(\mathfrak{a}))$ for all $\mathfrak{a} \in \mathcal{A}^{\overline{\mathfrak{p}}}$.

We in sum obtain

$$g(\overline{\mathfrak{p}}(\cdot),\cdot) \succsim g(\mathfrak{p}'(\cdot),\cdot) \sim \sum_{\mathfrak{a}\in\mathcal{A}} \overline{p}(\mathfrak{a}) g(\mathfrak{a}(\cdot),\cdot) \succsim \sum_{\mathfrak{a}\in\mathcal{A}} p'(\mathfrak{a}) g(\mathfrak{a}(\cdot),\cdot) \sim g(\overline{\mathfrak{p}}(\cdot),\cdot)$$

so that $\overline{\mathfrak{p}}$ and $\overline{p}$ are indeed indifferent. $\square$

Theorem 1 shows that modelling the agent's choice set as $\Delta\mathcal{A}$ is not without loss of generality. The assumption that the agent may choose any lottery over actions upon learning his type hides an assumption of dynamic consistency. The model of the agent's randomization device as an explicit algebra excavates this assumption. Part b) of Theorem 1 shows that the choice set of a semi-consistent agent is a subset of $\Delta\mathcal{A}$ that depends on the agent's preference over pure acts $\mathcal{A}$. In contrast, part a) shows that modelling the agent's choice set as $\Delta\mathcal{A}$ is without loss of generality if the agent is consistent. Theorem 1 applies whether the agent learns his signal before or after he learns the outcome of the randomization device. In the sequel I assume that the sets of lotteries $\Delta\mathcal{A}$ and $\Delta\mathcal{A}^*$ respectively are the choice sets of consistent and semi-consistent agents. Dropping the explicitly modeled randomization devices, call the $\succsim$-best lottery in $\Delta\mathcal{A}$ ( $\Delta\mathcal{A}^*$) the **optimal choice** or a consistent (semi-consistent) agent.

# 5 Bayesian agents

The Ellsberg Paradox shows that dynamic consistency may clash with strong independence. The present section firstly shows that this conflict never arises for Bayesian agents. Theorem 2 then shows that the converse also holds: Any dynamically consistent agent with access to a strongly independent universal randomization device must be Bayesian.

To understand the relation between dynamic consistency and strong independence for Baysians, consider a Bayesian agent with prior $\pi$. For this agent the universal randomization device $\Sigma^r$ is weakly independent of all other events $\mathcal{P}$ if and only if $\pi(E \cap E')$ equals $\pi(E)\pi(E')$ for all $E \in \Sigma^r$, $E' \in \mathcal{P}$. So weak independence reduces to standard stochastic independence in the Bayesian case. The same holds for strong independence as the Bayesian's preference over $\mathcal{P}$-measurable acts do not change with learning any (non-null) event $E \in \Sigma^r$, if and only if $\pi(E') = \pi(E' \mid E)$ holds for all $E' \in \mathcal{P}$.[10] Since for all $E' \in \mathcal{P}$ and all non-null $E \in \Sigma^r$ $\pi(E' \mid E) = \pi(E')$ holds if and and only if $\pi(E \cap E') = \pi(E)\pi(E')$, a Bayesian agent is consistent if and only if he is semi-consistent. In the Bayesian case dynamic consistency and strong independence then prescribe the same conditional preferences for randomization events: While Bayesian updating defines consistent conditional preferences, the weak-independence of $\Sigma^r$ from $\mathcal{P}$ implies their independence according to the agent's prior $\pi$.

Theorem 2 shows that Baysian agents are the only dynamically consistent and ambiguity averse agents who consider $\Sigma^r$ strongly independent. So any dynamically consistent agent with an ambiguity averse preference $\succsim$ considers $\Sigma^r$ to be strongly independent of $\mathcal{P}$ if and only if he is Bayesian.

**Theorem 2** *Say a dynamically consistent agent has an ambiguity averse preference $\succsim$. If this agent considers the randomization device $\Sigma^r$ strongly independent of $\mathcal{P}$, then $\succsim$ satisfies the independence axiom.*

**Proof** Suppose the agent's choices in some problem $(\succsim, \mathcal{Q}, g)$ violate the

[10]While strong and weak independence both reduce to the standard independence condition if the agent is a Bayesian, neither one implies the other without the assumption of ambiguity neutrality.

independence axiom. In particular, suppose that $p^* \in \Delta\mathcal{A}$ is an optimal choice in $(\succsim, \mathcal{Q}, g)$ for the agent, even though some $\mathfrak{a}^* \in \mathcal{A}$ with $p^*(\mathfrak{a}^*) > 0$ is not. Say the agent strictly prefers the pure plan $\tilde{\mathfrak{a}}$ to $\mathfrak{a}^*$, so $g(\tilde{\mathfrak{a}}(\cdot), \cdot) \succ g(\mathfrak{a}^*(\cdot), \cdot)$.

Define a $\Sigma^r \cap \mathcal{Q}$ measurable plan $\mathfrak{p}^* : \Omega \to A$ such that $\pi^r(E^{\mathfrak{p}^*}(\mathfrak{a})) = p^*(\mathfrak{a})$ for all pure plans $\mathfrak{a} \in \mathcal{A}$. By Theorem 1 $\mathfrak{p}^*$ is an optimal plan for the (consistent) agent. It must in particular be at least as good as the pure plan $\tilde{\mathfrak{a}}$, so $g(\mathfrak{p}^*(\cdot), \cdot) \succsim g(\tilde{\mathfrak{a}}(\cdot), \cdot)$. The consistency of the agent then implies $g(\mathfrak{p}^*(\cdot), \cdot) \succsim^{\mathfrak{p}^*}_{E^{\mathfrak{p}^*}(\mathfrak{a}^*)} g(\tilde{\mathfrak{a}}(\cdot), \cdot)$. Since $\mathfrak{p}^*$ restricted to $E^{\mathfrak{p}^*}(\mathfrak{a}^*)$ equals $\mathfrak{a}^*$ we then have $g(\mathfrak{a}^*(\cdot), \cdot) \succsim^{\mathfrak{p}^*}_{E^{\mathfrak{p}^*}(\mathfrak{a}^*)} g(\tilde{\mathfrak{a}}(\cdot), \cdot)$. Since the agent considers the randomization device $\Sigma^r$ to be strongly independent of $\mathcal{P}$, we then obtain the contradiction $g(\mathfrak{a}^*(\cdot), \cdot) \succsim g(\tilde{\mathfrak{a}}(\cdot), \cdot)$. $\square$

The trade off between independence and dynamic consistency described in Theorem 2 is reminiscent to the trade off between symmetry and dynamic consistency described in Epstein and Seo [11] and [12]. They show that a dynamically consistent agent, who considers a series of random experiments to be symmetric, must be ambiguity neutral. While Epstein and Seo [11] and [12] suggest various compromises between symmetry and dynamic consistency to allow for ambiguity aversion, I here suggest to curtail dynamic consistency just enough to allow for ambiguity aversion and independent randomization to co-exist.

# 6 Consistency: Observational Equivalence

Theorem 2 implies that we cannot simultaneously assume dynamic consistency, a strongly independent randomization device, and (strict) ambiguity aversion. To make room for ambiguity aversion, we can either weaken dynamic consistency or accept that agents change their preferences upon learning independent events. If we follow the second path out of this dilemma, different ambiguity attitudes do not manifest themselves in the agents' choices: Theorem 3 shows that any choice of a consistent agent can be rationalized as the choice of a Bayesian agent. To state Theorem 3 say the **Bayesian version** of a decision problem replaces the preference $\succsim$ over utility valued acts with an expected utility representation $\pi$, keeping all else equal.

**Theorem 3** *Fix a choice problem $(\succsim, \mathcal{Q}, g)$. If $p^* \in \Delta\mathcal{A}$ is an optimal choice for a consistent agent, then $p^*$ is an optimal choice in a Bayesian version of the problem.*

**Proof** Since the partitions of information and of payoff relevant events, $\mathcal{Q}^P$ and $\mathcal{R}^P$, are both finite we can represent any act $\sum_{\mathfrak{a}\in\mathcal{A}} p(\mathfrak{a})g(\mathfrak{a}(\cdot),\cdot)$ as a vector

$$\Big(\sum_{\mathfrak{a}\in A} p(\mathfrak{a})g\big(\mathfrak{a}(\theta),\sigma\big)\Big)_{\theta\in\mathcal{Q}^P,\sigma\in\mathcal{R}^P} \in \mathbb{R}^m \text{ with } m\colon =\mid \mathcal{Q}^P \mid \times \mid \mathcal{R}^P \mid .$$

The convex (and compact) hull of all such vectors $\mathcal{C}\colon = \{\sum_{\mathfrak{a}\in\mathcal{A}} p(\mathfrak{a})g(\mathfrak{a}(\cdot),\cdot) : p \in \Delta\mathcal{A}\}$ represents the set of all lotteries $\Delta\mathcal{A}$.

Let $f^*\colon = \sum_{\mathfrak{a}\in\mathcal{A}} p^*(\mathfrak{a})g(\mathfrak{a}(\cdot),\cdot)$. Since the agent is dynamically consistent $p^*$ is the $\succsim$-best lottery in $\Delta\mathcal{A}$, so $f^*$ represents the $\succsim$-best act in $\mathcal{C}$. The set[11] $\{f \in \mathbb{R}^m\colon f \gg f^*\}$ is convex and has $f^*$ is on its boundary. The monotonicity of $\succsim$ implies that $\{f \in \mathbb{R}^m\colon f \gg f^*\}$ is disjoint from $\mathcal{C}$. Since $f^*$ is also on the boundary of $\mathcal{C}$ there exists a separating hyperplane $H\colon = \{f \in \mathbb{R}^m\colon f\pi = f^*\pi\}$ with $f'\pi \leq f^*\pi < f\pi$ for all $f' \in \mathcal{C}$ and all $f \gg f^*$. Since $H \cap \{f \in \mathbb{R}^m\colon f \gg f^*\} = \emptyset$, $\pi(\theta\cap\sigma) \geq 0$ holds for all $\theta \in \mathcal{Q}^P, \sigma \in \mathcal{R}^P$ and $\pi(\tilde{\theta}\cap\tilde{\sigma}) > 0$ for some $\tilde{\theta}\cap\tilde{\sigma} \in \mathcal{Q}^P \cap \mathcal{R}^P$. So $\pi$ can be normalized to be a $\mathcal{Q}\cap\mathcal{R}$-measurable probability. A Bayesian with prior $\pi$ assigns the expected utility $f^*\pi$ to the utility valued act $f \in \mathcal{C}$ and is maxmizing his expected utility by choosing $p^*$.

$\square$

Theorem 3 extends Raiffa's [32] critique of the Ellsberg paradox to general choice problems with learning: No outside observer is able to discern the ambiguity attitudes of consistent agents. The separating hyperplane argument in the proof of Theorem 3 is, of course, not new. Most recently Kuzmics [26] applied it to a wide range of decision problems involving ambiguity. Kuzmics [26] reviews the long pedigree of the argument. Other

[11] The notation $f \gg f'$ stands for $f(\omega) > f'(\omega)$ for all $\omega \in \Omega$.

responses to the Raiffa's [32] critique can be found in the representations by Seo [36], Saito [34] and Ke and Zhan [21]. While Seo's [36] represenation assumes that randomization need not provide hedging, Saito [34] parameterizes the degree to which hedging is valuable. Ke and Zhang [21] let the value of hedging depend on the agents "subjective timing" of his own randomization and nature's move.

# 7 Semi-Consistency: Observational Difference

The introductory discussion of the Ellsberg paradox shows that semi-consistent and consistent agents behave differently. Theorem 4 shows that such behavioral differences may arise even in the simplest models with informative signals. Following the second example in the introduction, Theorem 4 considers the case where ambiguity only enters via the agent's informative signals. Theorem 4 furthermore assumes preferences with a MMEU-representation, that permit consistent consequentialist conditional preferences for all signals in the information partition $\mathcal{Q}^P$.

**Theorem 4** *Fix a decision problem $(\succsim, \mathcal{Q}, g)$, such that $\succsim$ has a MMEU-representation, ambiguity enters only via the agent's information partition $\mathcal{Q}^P$ and such that there exists a set of consequentialist and consistent conditional preferences for all signals in the information partition $\mathcal{Q}^P$. Then the set of optimal choices of a semi-consistent agent in the decision problem $(\succsim, \mathcal{Q}, g)$ may be disjoint from the set of optimal choices in any Bayesian version $(\pi, \mathcal{Q}, g)$ of $(\succsim, \mathcal{Q}, g)$.*

**Proof** Define $(\succsim, \mathcal{Q}, g)$ such that $\mathcal{R}^P = \{\rho, \lambda\}$ is the partition of payoff relevant events, $\mathcal{Q}^P = \{L, M, R\}$ is the information partition and $A :=$ $\{l, m, r\}$ the set of actions. Say $\succsim$ has a MMEU-representation with the set of beliefs $C^*$ which assigns probability $\frac{1}{2}$ to the two payoff relevant events $\lambda$ and $\rho$. So $\pi(\lambda) = \pi(\rho) = \frac{1}{2}$ holds for each $\pi \in C^*$ and the agent is an expected utility maximizer with respect to the partition of payoff relevant events. Furthermore assume that

$$C^*\colon=\big\{\pi\colon \pi(\lambda\cap L)=.5-\alpha,\qquad \pi(\lambda\cap M)=\alpha,\\ \pi(\rho\cap M)=.3-\alpha,\qquad \pi(\rho\cap R)=.2+\alpha,\quad \alpha\in[.1,.2]\big\}.$$

The following two tables illustrate the agent's MMEU. The first table lists $\pi(\theta\cap\sigma)$ for all $(\theta,\sigma)\in\{L,M,R\}\times\{\rho,\lambda\}$ parametrized by $\alpha\in[.1,.2]$. The second table lists the agent's utility $g(a,\sigma)$ for all $(a,\sigma)\in\{l,m,r\}\times\{\lambda,\rho\}$.

| | $L$ | $M$ | $R$ |
|---|---|---|---|
| $\lambda$ | $.5-\alpha$ | $\alpha$ | $0$ |
| $\rho$ | $0$ | $.3-\alpha$ | $.2+\alpha$ |

the set of priors $C^*$

| | $l$ | $m$ | $r$ |
|---|---|---|---|
| $\lambda$ | 21 | 10 | 0 |
| $\rho$ | 0 | 10 | 21 |

utilities $g(a,\sigma)$

To see that the decision problem permits a set of consequentialist consistent conditional preferences for all signals, we need some further definitions. For any set of beliefs $C$, say $C\mid_{\mathcal{Q}}\colon=\{\pi\mid_{\mathcal{Q}}\colon \pi\in C\}$ is the set of all of $\mathcal{Q}$-marginals $\pi\mid_{\mathcal{Q}}$ of all priors in $C$ and (for each $\theta\in\mathcal{Q}^P$) say $C(\cdot\mid\theta)\colon=\{\pi(\cdot\mid\theta)\colon\pi\in C\}$ is the set of $\theta$-updates $\pi(\cdot\mid\theta)$ for all $\pi\in C$. Epstein and Schneider [13] showed that any MMEU decision problem permits a set of consequentialist consistent conditional preferences $\{\succsim_\theta\}_{\theta\in\mathcal{Q}^P}$ if the set of beliefs $C$ is **rectangular** with respect to $\mathcal{Q}^P$ in the sense that

$$C=\Big\{\pi:\pi(\cdot\mid\theta)\in C(\cdot\mid\theta)\text{ for all }\theta\in\mathcal{Q}^P\text{ and }\pi\mid_{\mathcal{Q}}\in C\mid_{\mathcal{Q}}\Big\}$$

To see that the present set of beliefs $C^*$ is rectangular note that it can be represented as

$$\Big\{\pi:\pi(\cdot\mid\theta)\in C^*(\cdot\mid\theta)\text{ for all }\theta\in\mathcal{Q}^P\text{ and }\pi\mid_{\mathcal{Q}}\in C^*\mid_{\mathcal{Q}}\Big\}$$

with

$$C^*(\lambda\mid L)=\{1\},\ C^*(\lambda\mid M)=\Big[\frac{1}{3},\frac{2}{3}\Big],\text{ and }C^*(\lambda\mid R)=\{0\}$$
$$C^*\mid_{\mathcal{Q}}=\Big\{\big(\pi(L),\pi(M),\pi(R)\big)=\big(.5-\alpha,.3,.2+\alpha\big)\colon\alpha\in[.1,.2]\Big\}.$$

Since $C^*(\lambda\mid L)=\{1\}$ and $C^*(\lambda\mid R)=\{0\}$ the agent knows whether $\rho$ or $\lambda$ holds if he obtains signal $L$ or $R$. If the agent receives the signal $L$ or $R$

he is then best off by respectively choosing $l$ and $r$. The following calculation shows that $m$ is the agent's unique best choice when he receives the signal $M$.

$$
\begin{aligned}
\max_{\mathfrak{a}\in\mathcal{A}} \min_{\pi\in C^*} \sum_{\theta\in\mathcal{Q}^P,\sigma\in\mathcal{R}^P} \pi(\theta\cap\sigma) g(\mathfrak{a}(\theta),\sigma) = \\
\max_{\mathfrak{a}(M)\in A} \min_{\alpha\in[.1,.2]} \Big( (.5-\alpha) g(l,\lambda) + \alpha g\big(\mathfrak{a}(M),\lambda\big) + \\
(.3-\alpha) g\big(\mathfrak{a}(M),\rho\big) + (.2+\alpha) g(r,\rho) \Big) = \\
.7\times 21 + .3\Big( \max_{\mathfrak{a}(M)\in A} \min_{\gamma\in[\frac{1}{3},\frac{2}{3}]} (\gamma g\big(\mathfrak{a}(M),\lambda\big) + (1-\gamma) g\big(\mathfrak{a}(M),\rho\big) \Big) = \\
.7\times 21 + .3\Big( \min_{\gamma\in[\frac{1}{3},\frac{2}{3}]} (\gamma g(m,\lambda) + (1-\gamma) g(m,\rho) \Big).
\end{aligned}
$$

The first equality follows from the definition of $C^*$ and the optimality of $l = \mathfrak{a}^*(L)$ and $r = \mathfrak{a}^*(R)$ in the events $L$ and $R$. The second equality recognizes that the agent obtains utility $21 = g(l,\lambda) = g(r,\rho)$ if either $L$ or $R$ occurs which happens with the objective probability $.7 = .5 - \alpha + .2 + \alpha$. With the complementary probability .3 the agent faces a basic Ellsberg urn type problem, where $\gamma = \frac{\alpha}{\alpha + .3 - \alpha}$. In this problem the agent is best off choosing $m$ which yields the same utility in the two payoff relevant events. So $\mathfrak{a}^*$ with $\mathfrak{a}^*(L) = l$, $\mathfrak{a}^*(M) = m$, and $\mathfrak{a}^*(R) = r$ is the pure plan that maximizes $\min_{\pi\in C^*} \sum_{\theta\in\mathcal{Q}^P,\sigma\in\mathcal{R}^P} \pi(\theta\cap\sigma) g(\mathfrak{a}(\theta),\sigma)$; $\mathfrak{a}^*$ is therefore also the unique semi-consistent optimal choice.

However, any Bayesian with the same utility $g : A \times \Omega \to \mathbb{R}$ strictly prefers either $\mathfrak{a}'$ or $\mathfrak{a}''$ with $\mathfrak{a}'(M) = l$, $\mathfrak{a}''(M) = r$, and $\mathfrak{a}'(\theta) = \mathfrak{a}''(\theta) = \mathfrak{a}^*(\theta)$ for $\theta \in \{L, R\}$ to $\mathfrak{a}^*$. □

Just as in the second introductory example, ambiguity enters the above decision problem only via the agent's signals. In the above example it is however in the best interest of the agent to learn his signals. With a probability .7 the agent learns the payoff relevant event. This information is precise enough for the agent to accept the ambiguity that arises when he obtains the signal $M$.

# 8 Conclusion

The present study stands at the crossroads of three concepts: dynamic consistency, independence, and ambiguity. An event is ambiguous for an agent if he does not know its probability. The agent is ambiguity neutral if he is indifferent over all objective mixtures over sets of indifferent acts. A dynamically consistent agent follows through with any ex-ante optimal plan. Finally an event $E$ is independent of some some algebra $\mathcal{P}$ if learning $E$ never overturns the agent's ex-ante preference over $\mathcal{P}$-measurable acts.

Independence and dynamic consistency are compatible for ambiguity averse agents. For (strictly) ambiguity averse agents the two principles conflict. The behavior of any dynamically consistent agent can moreover be replicated as the behavior of an expected utility maximizer with a suitably defined prior. To create room for ambiguity averse preferences that also manifest themselves in the agent's behavior I suggest the model of dynamic semi-consistency. This model blends dynamic consistency and independence to derive conditional preferences. A semi-consistent agent is dynamically consistent with respect to informative signals, he does, however, not update his preferences upon learning uninformative independent events. The resulting dynamic inconsistency is enough for the behavior of ambiguity neutral agents to differ from the behavior of agents that are not.

Importing the present insights into game theory we see that any equilibrium in mixed strategies of a game with dynamically consistent and ambiguity averse agents is also the equilibrium of a Baysian version of the original game. This equivalence between Baysian and ambiguity averse game theory affects the equilibrium notions by Klibanoff [23], Lo [28], [29], Bade [4], Lehrer [27], Riedel and Sass [33], and Azrieli and Teper [3]. To generate novel predictions, applied works have then indeed not used these equilibrium notions. Ellis [9] and Kellner and Le Quement [22], for example, assume inconsistent updating to respectively study information aggregation and cheap talk with ambiguity averse agents. Similarly, Bose and Renou's [5] model of mechanism design relies on dynamic inconsistencies to ensure that the set of implementable social choice functions increases with the agents' ambiguity aversion.[12]

# References


[1] Anscombe, F. and R. Aumann: "A Definition of Subjective Probability" *The Annals of Mathematical Statistics*, 34, (1963), pp. 199-205.

[2] Aumann, R.: "Subjectivity and Correlation in Randomized Strategies", *Journal of Mathematical Economics*, 1, (1974), pp. 67-96.

[3] Azrieli, Y. and R. Teper: "Uncertainty Aversion and Equilibrium Existence in Games with Incomplete Information", *Games and Economic Behavior*, 73, (2011), pp. 310-317.

[4] Bade, S.: "Ambiguous Act Equilibria", *Games and Economic Behavior*, 71, (2011), pp. 246-260.

[5] Bose, S. and L. Renou: "Mechanism Design with Ambigous Communication Devices", *Econometrica*, 82 (2014), pp. 1853-1872.

[6] Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio, "Uncertainty Averse Preferences", *Journal of Economic Theory*, 146, (2011), pp. 1275-1330.

[7] di Tillio, A., N. Kos, and M. Messner: "The Design of Ambiguous Mechanisms", *Review of Economic Studies*, 84, (2017) pp. 237-276.

[8] Eichberger, J. and D. Kelsey: "Non-Additive Beliefs and Strategic Equilibria", *Games and Economic Behavior*, 30, (2000), pp. 183-215.

[9] Ellis, A.: "Condorcet Meets Ellsberg", *Theoretical Economics*, 11, (2016), pp. 865-895.

[10] Epstein, L and Le Breton, M.,: "Dynamically Consistent Beliefs Must be Bayesian", *Journal of Economic Theory*, 61, (1993), pp. 1-22.


[12] In their paper on mechanism design, Di Tillio, Kos, and Messner, [7] skirt the issue of dynamic consistency and independence by assuming that agents do not mix.

[11] Epstein, L, and K. Seo: "Symmetry of Evidence without Evidence of Eymmetry", *Theoretical Economics*, 5, (2010), pp 313-368.

[12] Epstein, L, and K. Seo: "Symmetry or Dynamic Consistency?", *The B.E. Journal of Theoretical Economics*, 11, (2011), pp 1-14.

[13] Epstein, L. and M. Schneider: "Recursive Multiple-Priors", *Journal of Economic Theory*, 113, (2003), pp. 1-31.

[14] Ghirardato, P.: "Revisiting Savage in a Conditional World", *Economic Theory*, 20, (2002), pp. 83-92.

[15] Gilboa, I. and D. Schmeidler: "Maxmin Expected Utility with Non-Unique Prior", *Journal of Mathematical Economics*, 18, (1989), pp. 141-153.

[16] Gilboa, I. and D. Schmeidler: "Updating Ambiguous Beliefs", *Journal of Economic Theory*, 59, (1993), pp. 33-49.

[17] Hanany, E. and P. Klibanoff: "Updating Preferences with Multiple Priors", *Theoretical Economics*, 2, (2007), pp. 261-298.

[18] Hanany, E. and P. Klibanoff: "Updating Ambiguity Averse Preferences", *The B.E. Journal of Theoretical Economics*, 9, (Advances), (2009) article 37.

[19] Hanany, E., P. Klibanoff, and S. Mukerji: "Incomplete Information Games with Ambiguity Averse Players" *American Economic Journal: Microeconomics*, 12, (2020) pp. 135-87.

[20] Hill, B.: "Dynamic Consistency and Ambiguity: A Reappraisal" *Games and Economic Behavior*, 120, (2020), pp. 289-310.
HEC Paris, 2013.

[21] Ke, S. and Q. Zhang: "Randomization and Ambiguity Aversion", *Econometrica*, 88, (2020), pp. 1159-1195.

[22] Kellner, C. and M. Le Quement: "Endogenous Ambiguity in Cheap Talk" *Journal of Economic Theory* 173, (2018), pp 1-17.

[23] Klibanoff, P.: "Uncertainty, Decision and Normal Form Games", mimeo, Northwestern (1996).

[24] Klibanoff, P.: "Stochastically Independent Randomization and Uncertainty Aversion", *Economic Theory*, 18, (2001), pp. 605-620.

[25] Klibanoff, P., M. Marinacci, and S. Mukerji: "A Smooth Model of Decision Making under Ambiguity", *Econometrica*, 73, (2005), pp. 1849-1892.

[26] Kuzmics, C: "Abraham Wald's Complete Class Theorem and Knightian Uncertainty" *Games and Economic Behavior*, 104, (2017), pp. 666-673.

[27] Lehrer, E.: "Partially Specified Probabilities: Decisions and Games" *American Economic Journal: Microeconomics* 4, (2012), pp. 70-100.

[28] Lo, K.C.: "Equilibrium in Beliefs under Uncertainty", *Journal of Economic Theory*, 71, (1996), pp. 443-484.

[29] Lo, K.C.: "Extensive Form Games with Uncertainty Averse Players", *Games and Economic Behavior*, 28, (1999), pp. 256 - 270.

[30] Machina M., "Dynamic Consistency and Non-Expected Utility Models of Choice under Uncertainty", *Journal of Economic Literature*, 27 (1989), 1622- 1668.

[31] McClennen, E.: *Rationality and dynamic choice: foundational explorations* (1990), Cambridge University Press, Cambridge.

[32] Raiffa, H.: "Risk, Ambiguity, and the Savage Axioms: Comment", *The Quarterly Journal of Economics*, 75, (1961), pp. 690-694.

[33] Riedel, F. and L. Sass: "Ellsberg Games", *Theory and Decision*, 76, (2013), pp. 469-509.

[34] Saito, K.: "Preferences for Flexibility and Randomization under Uncertainty", *American Economic Review*, 105, (2015), pp. 1246–1271.

[35] Siniscalchi, M. "Dynamic Choice under Ambiguity", *Theoretical Economics*, 6, (2011), pp. 379-421.

[36] Seo, K.: “Ambiguity and Second-Order Belief”, *Econometrica*, 77, (2009), pp.1575–1605. Kyoungwon Seo